Cell Therapy & Engineering Connect
ISSN: 3105-3866
2026, Vol. 2, Cite as: doi:10.x/journal.x.x.x

SCIFINITI PUBLISHING

OPEN ACCESS

Review Article

AI-Enabled Generative Design of Immune Cells and Receptors for Programmable Immunity

Ola A Al-Ewaidat[1] and Moawiah M. Naffaa[2]*

1. Department of Internal Medicine, Stanford University School of Medicine, Palo Alto, CA, USA
2. Independent Researcher, Mountain View, CA 94040, USA

* Author to whom correspondence should be addressed.

Abstract

Recent advances in generative artificial intelligence have begun to redefine the practice of cell and immune engineering. By learning the statistical and structural grammar of biological systems, generative models can now design T-cell receptors, chimeric antigen receptors, and synthetic immune circuits that meet complex objectives of affinity, stability, and specificity. When integrated into automated design–build–test–learn pipelines, these models enable continuous cycles of hypothesis generation, experimental validation, and model refinement, creating a closed feedback loop between computation and biology. This review examines how AI-driven generative design is transforming immunoengineering across multiple scales, from molecular recognition to cellular phenotype and clinical translation. It discusses the foundational architectures that support generative modeling in biology, the emergence of adaptive biofoundries that link digital design to manufacturing, and the translational pathways through which programmable immune cells may enter clinical practice. The review also explores the ethical and regulatory dimensions of algorithmic biology, emphasizing the need for transparency, equitable access, and anticipatory governance. Together, these developments signal the rise of a new paradigm, programmable immunity, in which biological design, therapeutic discovery, and ethical responsibility evolve within a single, intelligent framework.

Keywords

generative immunoengineering; programmable immunity; cell therapy; T-cell receptor design; chimeric antigen receptor (CAR-T); synthetic biology; AI-driven biomanufacturing; design–build–test–learn (DBTL) loop; biofoundry automation; algorithmic ethics; translational immunology; adaptive regulation.

1. Introduction

The immune system represents one of the most intricate adaptive architectures in biology, capable of sensing, learning, and remembering through dynamic molecular and cellular computation. Each T-cell receptor, B-cell receptor, and antibody embodies a combinatorial experiment in recognition, shaped by stochastic recombination and refined by selection [1]. This distributed intelligence endows immunity with extraordinary specificity and plasticity, but it also renders therapeutic design immensely complex. Efforts to reprogram immune function, whether by engineering monoclonal antibodies, constructing chimeric antigen receptors (CARs), or modulating regulatory cell lineages, have historically depended on rational and empirical strategies [2]. Such methods rely on human intuition, incremental mutagenesis, and extensive screening to optimize binding, stability, and signaling properties. Despite notable clinical success, this heuristic paradigm is inherently constrained. The potential design space of immune receptors and cell states is vast, multidimensional, and only sparsely sampled by experimentation [3,4].

For most of computational biology's history, artificial intelligence has been used primarily as an analytical or predictive tool. Discriminative and regression-based models have been trained to answer questions of the form: *What does this sequence do?* or *What structure will this sequence adopt?* Such models map existing biological inputs to inferred properties, enabling annotation, classification, and prediction. Landmark successes such as structure prediction with AlphaFold exemplify this paradigm, in which AI acts as a powerful interpreter of biological data rather than a creator of new biological entities, representing the peak of predictive modeling in biology [5]. In this framework, biology is treated as an object to be explained.

Generative artificial intelligence introduces a fundamentally different epistemology. Instead of asking what a given biological sequence means, generative models ask what sequences could exist that satisfy specified functional or structural goals. This represents a shift from inference to invention: from analyzing biology to actively proposing new biological forms. The core question becomes not "What does this sequence do?" but "What sequence will perform this function?" In this sense, generative AI transforms biology from a descriptive science into a design science, in which learned statistical representations of evolution and biophysics are used to create novel molecules, receptors, and cellular programs that have never existed in nature, as demonstrated by generative protein language models and structure-generating systems [6,7]

In recent years, the emergence of generative artificial intelligence has introduced a qualitatively new mode of biological design. Unlike traditional predictive algorithms that

classify or score existing data, generative models learn the underlying probability distributions of biological sequences, structures, and phenotypes, enabling them to synthesize novel entities consistent with the statistical grammar of life [8,9]. The same class of transformer and diffusion architectures that transformed natural-language and image generation, exemplified by models such as ESM-2, ProteinMPNN, and RFdiffusion, has been adapted to protein science, capturing contextual and geometric dependencies across millions of sequences [10,11]. These advances allow algorithms to infer latent representations that link sequence to structure and structure to function, providing a foundation for *de novo* design rather than retrospective optimization.

When applied to immunology, such models give rise to what can be described as generative immunoengineering, the computational creation of immune molecules and cellular programs guided by learned representations of biological principles. Early demonstrations illustrate this emerging capability. *PhysicoGPTCR* integrates large-language modeling with physicochemical conditioning to generate T-cell receptor (TCR) sequences with specified antigen context [12]. ProteinMPNN-based approaches have also been applied in study-specific workflows for immune-complex and TCR–pMHC interface design, coupling sequence generation with structural constraints to propose antigen-specific receptors that remain within realistic conformational manifolds [6]. Beyond receptor design, multimodal generative frameworks are beginning to link transcriptomic, proteomic, and signaling data to the prediction and ultimately the design of cellular phenotypes. These developments suggest that immune cells themselves may soon be programmable entities whose functional repertoires can be expanded algorithmically [13,14].

This transition from rational design to generative creation alters not only the technical workflow but also the epistemological foundations of cell engineering. The conventional design–build–test–learn (DBTL) cycle, a foundational framework in synthetic biology that iteratively connects computational design, experimental construction, performance testing, and data-driven learning, has been limited by experimental throughput. It is now being supplanted by a form of closed-loop intelligence, in which generative models and automated experimentation operate within adaptive DBTL frameworks under human oversight. In this architecture, model-generated hypotheses iteratively guide empirical validation, while experimental outcomes continuously refine model priors, forming a self-improving computational–biological feedback system [15]. The laboratory becomes a feedback system, an adaptive interface between computation and biology. Such integration has already accelerated antibody discovery, improved TCR-peptide–HLA binding prediction and inspired automated screening pipelines driven by active learning [16]. In the longer term, the coupling between *in silico* generation and *in vitro* verification may lead to self-improving

biomanufacturing ecosystems capable of producing tailored immune therapies with unprecedented speed and precision [17].

The opportunities of this paradigm are accompanied by significant challenges. The immune repertoire data that underpin generative training remain incomplete and biased toward particular species, diseases, and sequencing modalities, which raises concerns about generalizability and off-target risk [18]. The interpretability of high-capacity models is limited, and the regulatory frameworks governing AI-designed biologics are still emerging. Furthermore, the capacity to generate vast numbers of synthetic receptor sequences demands robust ethical and biosafety oversight to prevent unintended immunogenicity or dual-use misuse [19]. Addressing these issues will require collaborative standards that integrate computational transparency, experimental reproducibility, and normative guidance from clinical immunology and bioethics.

This review examines the convergence of artificial intelligence and immune cell engineering, focusing on how generative algorithms are reshaping the design landscape of receptors, signaling modules, and cellular behaviors. We synthesize conceptual foundations, survey recent technological advances, and outline the translational and governance challenges that accompany this accelerating field. The goal is to articulate a coherent framework for AI-enabled generative design of immune cells and receptors for programmable immunity, a paradigm that transforms the immune system from a biological phenomenon into a programmable platform in which data, learning, and design converge to expand the grammar of therapeutic possibility **(Figure 1)**. In this context, programmable immunity denotes the AI-enabled capacity to computationally design and regulate immune functions with defined precision, spanning receptor–antigen recognition, intracellular signaling, and cellular state transitions. It envisions an intelligent interface in which generative algorithms, molecular data, and synthetic biology converge to engineer immune behaviors in silico and validate them experimentally. Conceptually, programmable immunity reframes the immune system as a reconfigurable information network—continuously learnable, optimizable, and expressible through the generative grammar of biology.

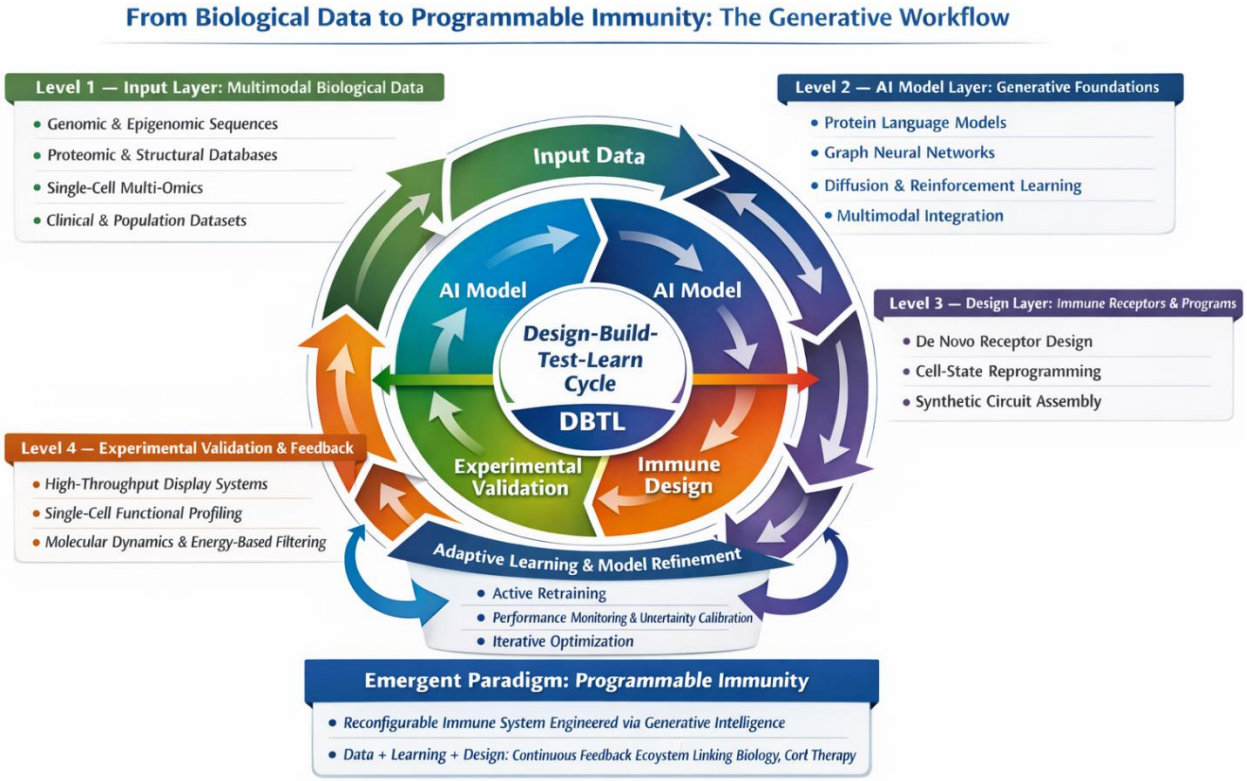

Figure 1. Foundations of Generative Immunoengineering.

2. The Generative Turn in Immunoengineering

For much of modern biotechnology, the design of immune therapeutics has followed a rational–empirical model rooted in target identification, scaffold selection, and iterative optimization. Breakthroughs such as monoclonal antibodies and the first generations of CAR-T cells emerged from this approach [20]. Yet the dependence on sequential mutagenesis and experimental screening imposes severe constraints on both scale and efficiency. The theoretical diversity of immune receptor sequences exceeds experimental capacity by many orders of magnitude [21], making comprehensive exploration of the antigen–receptor landscape unattainable and creating a persistent bottleneck in immune engineering [22].

The introduction of generative artificial intelligence has begun to transform this landscape. Unlike discriminative algorithms that categorize existing data, generative models learn the statistical patterns that define sequence–structure–function relationships and can therefore synthesize novel candidates consistent with those learned principles [23]. In this view, the immune system becomes not merely an object of analysis but a source of linguistic and structural priority from which algorithms infer the grammar of recognition [24]. Trained on large immune-repertoire datasets and receptor–antigen complexes, these models can

extrapolate to regions of sequence space that have not been sampled by natural evolution yet remain statistically and biophysically coherent. The result is a shift from selective discovery to probabilistic creation [25].

Generative design is inherently shaped by a tension between exploration and exploitation. On one hand, generative models are valued for their ability to explore unsampled regions of biological design space, proposing sequences and structures that have never existed in nature. On the other hand, their creativity is bounded by the data on which they are trained, leading to an inevitable bias toward exploiting learned statistical patterns. Excessive exploitation can cause generative systems to produce increasingly similar outputs over time, a phenomenon known as model collapse, in which diversity erodes and generated sequences become inbred reflections of the training set rather than genuinely novel designs.

This tension is particularly important in immunoengineering, where functional diversity is essential for discovering receptors with new specificity profiles. Models trained on biased or incomplete immune-repertoire datasets may overproduce familiar motifs while underexploring rare but therapeutically valuable configurations. Without explicit mechanisms to encourage diversity—such as entropy regularization, diversity-promoting sampling, adversarial training, or active-learning–driven exploration—generative systems risk converging on narrow regions of sequence space that appear safe but limit innovation [26,27].

Closed-loop design–build–test–learn frameworks partially mitigate this risk by reintroducing experimental feedback that expands training distributions beyond historical data. However, even in adaptive pipelines, careful balance is required between exploiting known high-performing designs and exploring uncertain regions that may harbor breakthrough solutions. Recognizing this exploration–exploitation tradeoff is therefore essential for responsible generative immunoengineering, tempering the optimism of algorithmic creativity with awareness of its statistical and epistemic limits.

Advances in protein foundation models provide computational architecture for this transformation. Transformer-based language models such as ESM-2, ProtT5, and MSA Transformer capture contextual dependencies across millions of sequences [28]. Diffusion and graph-neural architectures such as ProteinMPNN, RFdiffusion, and Chroma encode geometric constraints that preserve structural integrity [29]. When these models are fine-tuned on immunoglobulin, TCR, or antibody datasets, they learn both sequence syntax and physicochemical rules of folding and binding [30]. Conditional generation techniques allow the integration of antigenic or peptide–HLA information so that sequence generation

becomes *target-aware design* with affinity, specificity, and stability as explicit optimization goals [31].

Within immunoengineering, this generative capability is reshaping multiple domains. At the receptor level, models such as *PhysicoGPTCR* and *ProteinMPNN-TCR* co-model sequence, structure, and epitope context to generate receptor variants optimized for stability and antigen recognition [12,32]. Here, "ProteinMPNN-TCR" refers to study-specific application or fine-tuning of the open-source ProteinMPNN framework to TCR–pMHC or immune-interface design problems, rather than a distinct standalone software package [6]. At the construct level, emerging frameworks extend these generative principles to the modular design of chimeric receptors integrating variable fragments, hinge regions, and intracellular signaling domains to tune activation strength, expression, and safety [33]. Beyond receptors, multimodal generative models that integrate transcriptomic and proteomic profiles can infer regulatory or metabolic configurations that stabilize desirable cellular states. This emerging capability defines a new phase of AI-assisted cellular design [34].

The broader consequence of this transition lies in the restructuring of experimental reasoning. Traditional discovery pipelines rely on experimental data to generate hypotheses. In the generative framework, models propose candidates that guide experiments, and experimental outcomes continually refine model priors, creating a self-reinforcing design–build–test–learn cycle [17]. Laboratories are evolving into adaptive systems in which computational and biological processes operate in tandem, accelerating the translation of insight into function [35]. This feedback architecture mirrors closed-loop control in engineering and signals the rise of self-optimizing biomanufacturing platforms.

Despite these advances, the generative turn requires careful evaluation. Models trained in incomplete or biased repertoires may produce sequences that violate structural or safety constraints [36]. Ensuring safety, interpretability, and reproducibility in AI-generated biologics demands comprehensive benchmarking, transparent documentation, and the establishment of regulatory frameworks suited to algorithmic design [37]. Generative systems should therefore be viewed as collaborators rather than replacements for scientific expertise, augmenting experimental insight rather than automating it.

Viewed through this lens, generative immunoengineering is both a technological and conceptual redefinition. Computation no longer merely represents biological systems; it participates in their creation. The ability to generate plausible, functional immune receptors *ab initio* transforms design from an empirical craft into an algorithmic discipline and expands the creative frontier of synthetic biology [38]. The subsequent sections examine the architecture, data resources, and experimental integrations that together define this

emerging paradigm. A convergence of artificial intelligence and immunological engineering directed toward programmable immunity.

3. Foundations of Generative Biology for Immune Systems

The application of generative artificial intelligence to immunoengineering rests upon foundational advances in computational biology that have redefined how protein sequences, structures, and functions are represented. Over the past five years, models originally designed for natural language processing have been adapted to the protein domain, where amino acid sequences are treated as biological sentences governed by an implicit grammar of evolution [39]. This conceptual analogy between linguistic and molecular syntax has allowed transformer architectures, recurrent neural networks, and diffusion-based frameworks to learn the contextual dependencies that link sequence motifs to structural and functional outcomes. These models now provide the representational backbone for generative biology, enabling the creation of new molecular entities that adhere to the learned statistical rules of natural proteins [40].

At the core of this transformation lies the protein language model (pLM). Early models such as UniRep and ProtBERT demonstrated that contextual embeddings derived from millions of sequences could capture latent biophysical properties including secondary structure propensity, binding-site probability, and thermostability [41,42]. More recent architectures such as ESM-2 and MSA Transformer incorporate evolutionary and multiple-sequence alignment information, achieving state-of-the-art performance in structure prediction and mutational effect inference [43,44,45]. These embeddings serve as a universal representation that can be fine-tuned for diverse downstream tasks, including receptor generation, antigen classification, and affinity optimization. In the context of immune receptor design, such representations provide a high-dimensional landscape in which antigen specificity and receptor stability can be jointly optimized by generative sampling [46].

These high-dimensional embeddings define what can be viewed as a latent space: a continuous representational manifold in which biological properties such as stability, affinity, specificity, and phenotype are smoothly encoded. Rather than serving only as an internal technical feature of neural networks, this latent space functions as a new design canvas for biologists. Each point in this space corresponds to a potential biological design, and nearby points represent variants that differ gradually in structure or function. In this sense, generative biology operates not by discrete trial-and-error but by navigation within a learned biological landscape [7,42].

Generative design occurs through steering within this latent space. Researchers do not manipulate sequences directly; instead, they control sampling by adjusting conditioning variables such as antigen identity, structural templates, physicochemical constraints, or desired cellular states. These conditioning variables bias the regions of latent space from which designs are drawn, effectively allowing researchers to "navigate" toward functionally meaningful zones. Interpolation within latent space enables the generation of intermediate designs that smoothly trade off between properties, such as affinity versus stability or activation versus persistence [7,47].

However, the geometry of this design canvas is entirely determined by the data used to construct it. Latent spaces inherit the biases, omissions, and distortions of their training datasets. If immune-repertoire or structural datasets overrepresent particular species, diseases, or experimental systems, the resulting latent space will privilege those regions of biology while neglecting others. As a result, generative models may produce designs that are biologically plausible yet therapeutically irrelevant or unsafe, because they are optimized within a distorted representation of biological reality [48].

In generative immunoengineering, data quality dictates design destiny. The clinical relevance, novelty, and safety of generated receptors or cell states cannot exceed the informational boundaries of the data from which their latent spaces are learned. Curated diversity, balanced representation, and deep functional annotation are therefore not auxiliary concerns; they are the primary determinants of whether generative models yield meaningful therapies or merely elegant but clinically empty designs [27].

Complementary to language-based methods are structure-aware generative models that explicitly incorporate geometric and energetic constraints. Graph neural networks, energy-based models, and diffusion frameworks such as ProteinMPNN, RFdiffusion, and Chroma model the conditional probability of atomic arrangements given a target fold or binding interface [29,49,50]. By learning from experimentally determined structures and molecular dynamics simulations, these models capture the geometric invariants that define stable tertiary and quaternary conformations. When applied to immune complexes, structure-aware generators can propose receptor sequences that maintain structural fidelity while accommodating specific epitope geometries, a key requirement for accurate antigen engagement[17]. The ability to jointly model sequence and structure distinguish these approaches from earlier heuristic design pipelines and provides a direct pathway from digital generation to physical synthesis.

3.1. Technical Adaptation of Generative Architectures for Immunology

Although transformer and diffusion architectures originate from general-purpose machine learning, their application to immunology requires domain-specific adaptation of inputs, conditioning strategies, and physical constraints. In immune receptor design, the biological meaning of sequence context, evolutionary history, and structural feasibility must be explicitly encoded rather than treated as abstract tokens.

Multiple sequence alignments (MSAs) are processed in protein language models by representing aligned residues across homologous sequences as parallel input channels rather than linear strings. In models such as the MSA Transformer, each alignment column is embedded jointly across sequences, allowing attention mechanisms to learn evolutionary covariation between residues [5,51]. For immune receptors, however, MSAs are often sparse or biased, particularly for TCRs, where extreme diversity and limited structural sampling reduce alignment depth. As a result, immune-focused adaptations either restrict MSA usage to framework regions, augment alignments with synthetic or inferred homologs, or replace MSAs with large-scale repertoire embeddings learned directly from unaligned sequences. In this context, models such as ESM-2 rely primarily on self-attention over raw sequences [44], allowing them to infer contextual constraints without requiring dense evolutionary alignments.

Physical constraints are incorporated through architecture design and post-generation filtering. Structure-aware models such as ProteinMPNN and RFdiffusion encode proteins as graphs or geometric point clouds, where residues are nodes and spatial relationships define edges [6,52,53]. These models learn conditional probability distributions over amino-acid identities or atomic coordinates given a fixed backbone, interface, or motif. In diffusion models, generation proceeds through iterative denoising steps that progressively enforce geometric consistency, steric feasibility, and backbone continuity. Training on experimentally determined structures allows these models to internalize folding rules, hydrogen-bonding patterns, and packing constraints implicitly through loss functions that penalize geometric deviation.

In immune applications, these physical constraints are further specialized by conditioning on antigenic context. For TCR design, structural templates of TCR–pMHC complexes define spatial constraints on complementarity-determining regions, particularly the CDR3 loop, which exhibits high flexibility [54,55]. Diffusion-based and structure-conditioned models handle this extreme flexibility by decoupling rigid and flexible regions during generation. Framework regions are typically fixed or tightly constrained using backbone templates derived from known TCR structures, preserving global fold stability. In contrast, the CDR3

loop is either partially masked or assigned higher stochastic freedom during diffusion, allowing its backbone and side-chain geometry to be resampled more extensively. Conditional variables such as interface residue masks or antigen-contact maps restrict this flexibility to antigen-facing regions, ensuring that variability is focused on functional interface residues rather than destabilizing the overall receptor. This region-specific flexibility enables diffusion models to explore diverse CDR3 conformations while maintaining structural integrity of the conserved TCR scaffold. Conditioning variables may include peptide–HLA embeddings, interface residue masks, or physicochemical features that bias generation toward shapes compatible with specific epitopes.

In addition to architecture-level constraints, most pipelines integrate physics-based post-processing. Generated sequences or structures are filtered using molecular-dynamics relaxation, energy minimization, or docking simulations to eliminate candidates that violate thermodynamic stability or steric feasibility [56,57]. This hybrid approach ensures that statistical plausibility learned by neural networks is grounded in physical reality, reducing the risk of generating biologically impossible but mathematically plausible designs.

To clarify how these architectures differ in data requirements, modeling strategy, accessibility, and immune-engineering use cases, a comparative overview of major generative frameworks used in immunoengineering is provided in supplementary **Table S1**.

While multiple generative architectures are now used in immunoengineering, they differ fundamentally in the type of biological information they operate on. Protein language models such as ESM-2 primarily learn from sequence statistics and are well suited for exploring diversity in TCR or antibody repertoires, whereas structure-aware models such as ProteinMPNN and RFdiffusion incorporate geometric constraints and are better suited for interface design and de novo binder generation. TCR-focused frameworks such as PhysicoGPTCR and ProteinMPNN-TCR represent emerging efforts to adapt these general-purpose architectures to immune-specific constraints, particularly antigen context and complementarity-determining region variability. In contrast, CAR-T design has largely relied on generative or reinforcement-learning frameworks trained on modular construct libraries, reflecting the engineering rather than evolutionary nature of CAR architectures.

3.2. Code Availability and Reproducibility

Given the rapid evolution of generative immunoengineering, transparency and code accessibility are essential for reproducibility and independent benchmarking. The generative frameworks discussed in this review span open-source tools, research prototypes, and proprietary or study-specific systems.

Several widely used foundation models are openly available. ESM-2 and related protein language models are released by Meta AI with open-source code and pretrained weights [44]. ProteinMPNN, RFdiffusion, Chroma, and MSA Transformer are also distributed as open-source research tools with publicly accessible repositories, enabling independent replication and extension of published results [6,53,58].

In contrast, a number of immune-specialized frameworks such as PhysicoGPTCR and ProteinMPNN-TCR represent research prototypes whose code availability varies by study [59]. Some implementations are partially released or shared upon request, while others remain internal to the originating research groups. Where code is not publicly available, we explicitly label these frameworks as research prototypes rather than community tools.

Generative CAR-design frameworks are frequently developed within academic–industrial collaborations or commercial platforms and are often proprietary or study-specific [60,61]. In such cases, methodological details are available through publications, but source code and trained models are not released.

Each framework discussed is explicitly labeled as open-source, partially open, or proprietary **(Supplementary Table S1)**, and GitHub links are provided for tools with public repositories. Frameworks that are conceptual, hypothetical, or not accompanied by released code are clearly identified as such to avoid overstating their reproducibility.

Recent innovations have begun to integrate multimodal generative architectures that combine sequence, structure, and system-level data into unified frameworks. Variational autoencoders and diffusion models trained on multi-omic datasets can embed gene expression, protein abundance, and signaling dynamics within shared latent spaces [62,63]. These models enable not only receptor-level design but also the generation of synthetic cell states, offering predictions of how genetic or metabolic perturbations may influence functional phenotypes. When aligned with single-cell RNA sequencing, proteomic profiling, and CRISPR perturbation data, multimodal models allow researchers to simulate the outcomes of cellular reprogramming before executing experimental interventions [64]. This integration forms the conceptual foundation of generative immunoengineering, where molecular design and cell-state control converge through shared representational learning.

Another essential component of this foundation is conditional generation, which introduces explicit control variables into the generative process. Conditioning can be achieved through structural templates, physicochemical features, antigenic context, or phenotype-level objectives [65]. In receptor engineering, conditional models can generate sequences that maximize predicted binding affinity to a specified epitope while minimizing cross-reactivity and immunogenicity [66,67]. In cell engineering, conditioning may involve the specification

of transcriptional or metabolic profiles corresponding to desired states such as resistance to exhaustion, enhanced memory formation, or altered cytokine secretion [68,69]. By encoding these objectives into the generative process, models move beyond unsupervised creativity toward constrained biological design that aligns with therapeutic goals.

The reliability of generative models depends critically on the quality, diversity, and annotation of training data. Immune-repertoire datasets such as IEDB, VDJdb, OAS, and PIRD provide millions of receptor sequences, yet these datasets are biased toward particular species, disease contexts, and sequencing platforms [70]. Structural datasets remain comparatively sparse, limiting model generalization for certain receptor classes or antigen types. Integrating curated experimental datasets with synthetic augmentation strategies, including contrastive learning and adversarial perturbation, has emerged as a strategy to expand functional diversity while maintaining biological plausibility [71]. Standardization of data formats and metadata annotations is equally important for ensuring reproducibility across laboratories and for benchmarking model performance under transparent conditions [72].

The theoretical strength of these models is matched by their capacity for interpretability and embedding analysis, which enables a mechanistic understanding of learned representations. Techniques such as attention visualization, feature attribution, and latent-space interpolation have revealed that protein language models implicitly capture biochemical hierarchies reminiscent of evolutionary phylogenies [43,73]. In immune systems, these embeddings encode information about complementarity-determining region composition, binding topology, and germline lineage relationships [74]. Understanding how models internalize these features not only enhances trust in generative predictions but also provides a new lens through which to examine the informational logic of the immune repertoire itself.

Together, these methodological pillars establish the computational grammar of generative biology. The convergence of language-based, structure-aware, and multimodal approaches provides the mathematical substrate on which immune receptor and cell-state generation can occur with controllable precision. As models become larger and more contextually integrated, they begin to approximate a universal generator of biomolecular function, one capable of producing sequences, folds, and regulatory motifs that are both novel and biologically viable. In the context of immunoengineering, this synthesis enables the design of receptors, signaling networks, and phenotypic programs guided not by human intuition alone but by statistical representations of evolution and function embedded within artificial intelligence [75,76].

The next section will examine how these generative foundations are being applied to the design of immune repertoires and antigen-specific recognition, focusing on model architectures, conditioning strategies, and validation frameworks that connect digital generation to experimental reality.

4. Learning Immune Specificity from Repertoires and Structures

The capacity of the adaptive immune system to distinguish self from non-self and to recognize an effectively unlimited range of antigens arises from the extraordinary diversity of its receptor repertoires. Each T-cell and B-cell receptor represents a molecular hypothesis drawn from the combinatorial space of V(D)J recombination, junctional variability, and somatic hypermutation [77]. Mapping this diversity and translating it into actionable design principles have long challenged immunology. Generative artificial intelligence now provides new tools for capturing the probabilistic architecture of immune specificity by learning directly from large-scale receptor and antigen datasets [78,79].

Immune-repertoire learning relies on sequence datasets such as IEDB, VDJdb, OAS, and PIRD, which together contain millions of annotated T-cell and B-cell receptor sequences linked to antigenic or disease contexts [80,81,82]. Deep representation models trained on these corpora can learn statistical signatures that define clonotype structure, CDR usage, and gene-segment pairing preferences [83]. Early models such as DeepTCR, TCR-BERT, and Immune2Vec demonstrated that unsupervised embeddings derived from raw sequence data capture functional and evolutionary relationships between receptors [4,84,85]. These embeddings have subsequently become the basis for generative modeling, enabling conditional sampling of sequences that preserve repertoire-level statistics while exploring unsampled regions of sequence space.

The extension of this approach to structure-informed modeling has further refined our understanding of immune recognition. High-resolution structural data from crystallography and cryo-electron microscopy, combined with molecular-dynamics simulations, provide explicit information about how complementarity-determining loops engage peptide–MHC complexes or conformational epitopes [86]. Graph neural networks and diffusion-based models such as ProteinMPNN-TCR, AlphaBind, and ImmuneDiffusion encode both sequential and geometric features, allowing generation or evaluation of receptor variants that maintain structural stability while optimizing epitope complementarity [66,87,88]. By unifying sequence and structural representations, these models can predict or design receptors that balance affinity, cross-reactivity, and biophysical feasibility.

4.1. Data Scarcity and Strategies for Learning Under Limited Structural Supervision

A fundamental bottleneck in TCR-specific generative design is the scarcity of high-quality paired TCR–pMHC structural data [59,89]. Compared with general protein databases, experimentally resolved immune complexes represent only a small and biased subset of possible receptor–antigen interactions. This limitation constrains direct supervised training of structure-aware generative models and necessitates strategies that learn under weak or indirect supervision.

One major strategy is transfer learning from large, non-immune protein corpora. Foundation models such as ESM-2, ProteinMPNN, and RFdiffusion are first trained on millions of generic protein sequences or structures, allowing them to internalize universal rules of folding, packing, and interface geometry [6,44,53,90]. Immune-specific tasks then rely on fine-tuning or conditioning using relatively small TCR–pMHC datasets, effectively leveraging general protein knowledge to compensate for immune-data sparsity.

A second strategy involves using large-scale immune-repertoire sequencing as a proxy for structural supervision [4,91]. Although most repertoire datasets lack paired antigen or structure information, they capture statistical regularities of CDR usage, V(D)J recombination, and junctional diversity. Generative models trained on these repertoires learn the grammatical constraints of immune receptors, which can then be coupled to structural or antigen-conditioned models through multimodal or hierarchical architectures.

Synthetic data augmentation provides an additional route to overcoming small-data limitations. Structure prediction tools, docking algorithms, and molecular-dynamics simulations are used to generate approximate TCR–pMHC complexes from known sequences, expanding training sets beyond experimentally resolved structures [92,93]. Although these synthetic structures carry uncertainty, they enable models to explore broader interface geometries and can be filtered using energy-based or stability criteria.

Active learning and closed-loop DBTL frameworks further mitigate data scarcity [94,95]. Generative models propose receptor candidates whose experimental testing is expected to yield maximal information gain. High-throughput binding assays, display systems, and single-cell functional screens then generate new labeled data, which are fed back into model retraining. This iterative strategy allows models to progressively improve even when initial labeled datasets are small.

Finally, weakly supervised and contrastive-learning approaches allow models to learn from partially labeled data, such as receptors known to bind a class of antigens without precise

structural resolution [96,97]. By learning relative similarities rather than absolute labels, these methods extract useful signal from noisy or incomplete datasets and reduce dependence on fully resolved structures.

Learning immune specificity also depends on understanding contextual conditioning, in which receptor generation is guided by features of the antigen or by the cellular environment in which binding occurs. Conditional generative frameworks incorporate peptide–HLA embeddings, physicochemical descriptors, and even transcriptomic signatures of the responding cell population [98]. This conditioning allows models to generate receptors tailored to particular epitopes or immunological niches rather than relying on global repertoire statistics. The resulting designs can be filtered by computational docking, binding-energy prediction, or molecular-dynamics relaxation to ensure structural plausibility and to screen for potential off-target interactions [99,100].

An equally important development is the use of contrastive and active-learning strategies that couple model training to experimental feedback. High-throughput binding assays, yeast or mammalian display systems, and single-cell sequencing technologies provide empirical data that refine model priors through iterative updates [101,102,103]. Active learning enables the model to identify regions of uncertainty and to propose new receptor sequences whose experimental testing would maximize information gain. This closed-loop process progressively enhances model fidelity and creates an adaptive design framework that mirrors the iterative nature of immune evolution itself.

Despite these advances, challenges remain in ensuring that learned representations reflect biological causality rather than statistical correlation. Sequence redundancy, sampling bias, and limited negative examples can inflate apparent model accuracy while masking gaps in functional understanding [104]. The field is responding through the creation of standardized benchmarking platforms, curated cross-reactivity datasets, and transparent reporting of validation metrics that measure generalization across antigen classes and experimental systems [105,106]. Incorporating structural energetics, thermodynamic parameters, and molecular-simulation outputs into the training regime further grounds model predictions in biophysical reality [107,108].

Together, these developments illustrate a convergence between data-driven learning and structural immunology. By jointly modeling sequence, structure, and antigenic context, generative and representation models are beginning to reconstruct the rules that govern immune recognition. In practical terms, they offer the capacity to generate receptor repertoires that are both diverse and functionally directed, to predict the cross-reactivity landscape of candidate therapeutics, and to guide the rational expansion of immune

libraries toward desired antigen spaces. The synthesis of repertoire-scale learning with structural modeling thus forms the methodological core of generative immunoengineering, providing the analytical framework through which specificity can be understood, predicted, and designed **(Supplementary Table S2)**.

5. AI-Driven Design of Immune Receptors and Constructs

The emergence of generative artificial intelligence has transformed the design of immune receptors from an empirical pursuit into a computational discipline. Models that learn sequence–structure–function relationships across vast biological corpora can now propose receptor variants and modular constructs that meet predefined design constraints for affinity, stability, signaling balance, and manufacturability [109]. This integration of algorithmic inference with molecular immunology redefines the creative boundary of immune engineering, transforming receptor discovery into a process of directed generation guided by learned biological priors [110].

5.1 Designing Antigen-Specific Receptors

The adaptive immune system recognizes antigens through an immense repertoire of receptor sequences that encode highly specific binding topologies. Generative models trained on large-scale T-cell receptor (TCR) and antibody repertoires can recapitulate and extend this natural diversity. Transformer-based language models, such as PhysicoGPTCR (a physics-informed generative framework for TCR design), employ contextual embeddings that capture residue-level physicochemical properties while conditioning generation on peptide–HLA or epitope descriptors [12,111,112]. By sampling from latent spaces that integrate both sequence statistics and antigenic context, these models generate plausible receptor candidates that occupy unobserved yet biologically coherent regions of sequence space. A schematic overview of the PhysicoGPTCR architecture and its conditioning,generation, and filtering stages is shown in **Figure 2**.

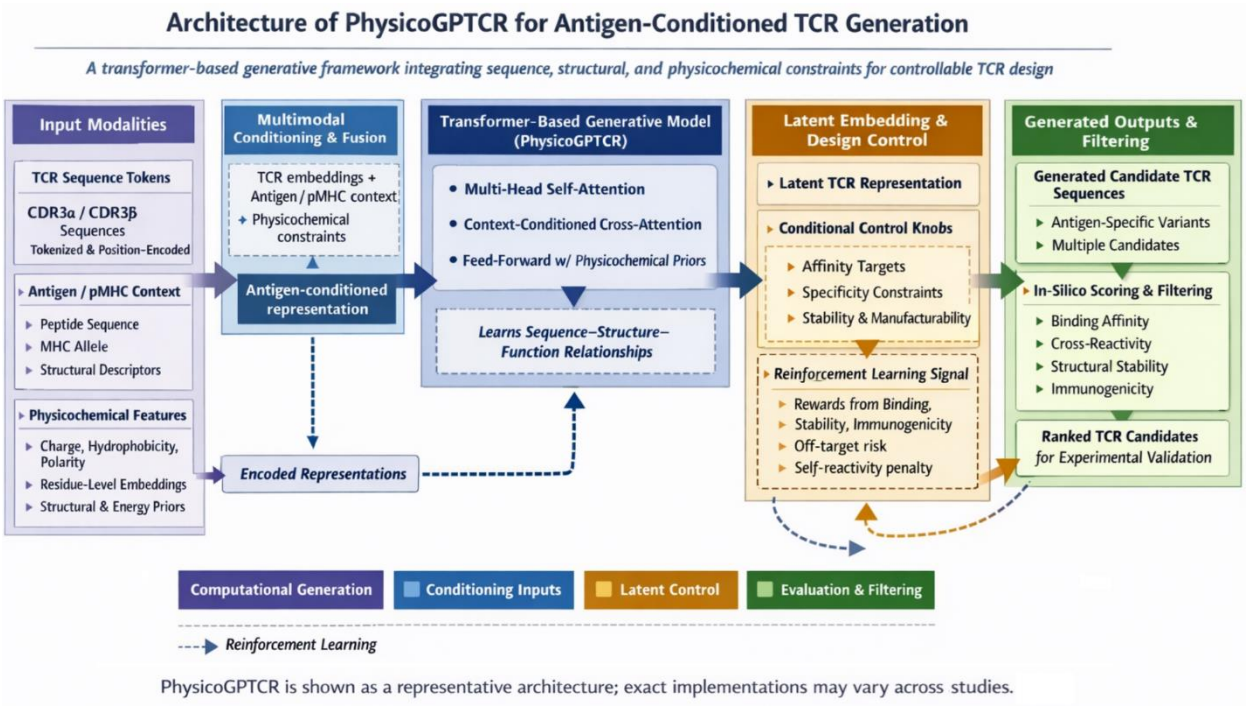

Figure 2. Architecture of PhysicoGPTCR for antigen-conditioned TCR generation.

In practice, these objectives are rarely independent and often compete. For example, increasing receptor affinity can compromise structural stability, broaden cross-reactivity, or elevate immunogenicity. Generative design therefore operates as a multi-objective optimization problem rather than a search for a single optimal solution. Reinforcement-learning–based and active-learning frameworks address this challenge by optimizing over multiple reward components simultaneously, identifying Pareto-optimal fronts that represent families of non-dominated receptor designs balancing competing objectives. Instead of collapsing design into a single scalar score, these models preserve trade-offs between affinity, specificity, stability, and safety, allowing human designers to select candidates that align with therapeutic priorities. In this sense, human-defined reward functions and constraints explicitly shape the evolutionary trajectory of AI-designed receptors, guiding exploration toward clinically acceptable regions of design space rather than maximal affinity alone [113,114,115,116,117].

5.1.1 Experimental Validation of AI-Designed TCRs

While many generative frameworks remain at the proof-of-concept stage, a growing number of studies have demonstrated wet-lab validation of AI-designed or AI-optimized TCRs. These validations typically involve surface expression analysis by flow cytometry, antigen-specific binding assays, cytokine release, and target-cell killing assays.

In one representative study, antigen-conditioned generative models were used to design novel TCR sequences targeting defined peptide–HLA complexes. Selected candidates were expressed in primary T cells and showed specific peptide–HLA binding by flow cytometry and multimer staining, confirming that the generated sequences formed functional surface receptors [67,118,119].

Functional validation further demonstrated that AI-designed TCR-T cells could mediate antigen-specific activation. Upon co-culture with antigen-positive target cells, engineered T cells exhibited increased CD69 or CD25 expression, secreted interferon-γ or IL-2, and selectively lysed antigen-positive but not antigen-negative targets in cytotoxicity assays [67,120].

In another example, structure-guided or diffusion-based design was used to generate TCR variants predicted to improve interface complementarity with tumor antigens. Experimental testing showed enhanced binding affinity relative to parental receptors, accompanied by increased killing efficiency in chromium-release or luminescence-based cytotoxicity assays [120,121,122].

Together, these studies demonstrate that generative models are not merely theoretical design tools but can produce receptors that function in living T cells, supporting antigen recognition, signaling, and target-cell killing under experimental conditions.

5.1.2 Safety, Cross-Reactivity, and Off-Target Risk Mitigation

In immunotherapy, design without safety is clinically meaningless. High-affinity receptors that cross-react with healthy tissues have caused fatal toxicities in multiple clinical trials, underscoring that specificity and safety must be co-optimized rather than treated as secondary objectives [123,124].

Generative models address cross-reactivity by incorporating negative design constraints. Instead of optimizing only for binding to a target peptide–HLA complex, models are trained or filtered against large libraries of self-peptides, tissue-specific antigens, and predicted off-target epitopes [121,125]. Candidate receptors are scored not only by predicted affinity for the intended antigen but also by similarity of their binding interfaces to receptors known to recognize self-antigens. This multi-objective optimization penalizes designs that exhibit high predicted binding to unintended targets.

Sequence-level safety screening is performed using immunogenicity and humanization predictors. Generated receptors are evaluated for similarity to germline human TCR frameworks, for the presence of rare or foreign motifs, and for predicted immunogenic epitopes. Models trained on large human repertoires implicitly bias generation toward

human-like sequences, while post-generation filters remove candidates with high predicted risk of host immune recognition. In practice, this includes the use of T-cell epitope–prediction tools and MHC-binding predictors such as NetMHCpan-4.0/4.1 to estimate peptide presentation likelihood and potential immunogenicity, combined with deep-learning immunogenicity scoring frameworks such as DeepImmuno that evaluate peptide–MHC immunogenic potential beyond simple binding affinity [126,127]. These predictors act as negative-selection layers, ensuring that high-affinity designs are not advanced if they carry elevated risk of host immune responses. Models trained on human repertoires implicitly bias generation toward human-like sequences, while post-generation filters remove candidates with high predicted risk of host immune recognition.

Structure-aware and diffusion-based models further enable geometric screening of cross-reactivity. By modeling receptor–antigen interfaces in three dimensions, these frameworks can dock generated receptors against panels of self-peptides and structurally related epitopes [53,128]. Candidates that form stable interfaces with non-target peptides are eliminated through energy-based or steric-clash criteria, reducing the risk of unintended tissue targeting.

Safety is further reinforced through physics-based and experimental filtering. Generated receptors are subjected to molecular-dynamics relaxation, energy minimization, and docking simulations to identify unstable or promiscuous interfaces. In experimental pipelines, early-stage screening includes testing against panels of healthy-cell antigens using flow cytometry, multimer staining, and co-culture assays to detect unintended activation before therapeutic advancement [129,130].

Together, these strategies transform generative design into safety-aware design, in which affinity, specificity, stability, and off-target risk are optimized simultaneously. In this paradigm, generative models do not merely create receptors but act as risk-filtering systems that integrate statistical learning, structural modeling, and biological priors to minimize the probability of lethal cross-reactivity.

5.1.3 Evaluation Metrics for Generative TCR Design

Evaluation metrics for generative TCR design span four interconnected levels: structural plausibility, biophysical stability, antigen-binding performance, and cellular functional output. The performance of generative models in TCR engineering cannot be judged by sequence novelty alone. Instead, evaluation requires a multilevel framework that integrates structural confidence, biophysical stability, antigen-binding strength, specificity, and functional cellular output.

At the structural level, predicted folding confidence is commonly assessed using metrics such as predicted Local Distance Difference Test (pLDDT) scores from structure-prediction models and root-mean-square deviation (RMSD) relative to known or docked templates [5,131]. High pLDDT values and low RMSD indicate that generated receptors adopt physically plausible conformations compatible with stable expression and surface display.

When benchmarking interface accuracy for TCR–pMHC complexes, it is important to distinguish complex-structure prediction (e.g., homology modeling and AlphaFold-Multimer–style "fold-and-dock") from generative design (sequence/interface generation). For structure prediction, recent assessments show that AlphaFold-derived pipelines can yield near-native TCR–pMHC geometries in a subset of cases, but performance remains variable and is often limited by docking orientation errors and by the flexibility of CDR3 loops; therefore, interface-focused metrics such as iRMSD, DockQ, and contact recovery are frequently reported alongside global RMSD [128,132]. For generative design, most studies do not yet report standardized head-to-head comparisons to homology modeling or AlphaFold-Multimer on the same benchmark sets; instead, structure-aware generative pipelines typically use AlphaFold/AlphaFold-Multimer (or related predictors) as a screening and validation layer, reporting predicted interface RMSD/PAE proxies, docking scores, and then prioritizing candidates for experimental binding and functional testing. We therefore summarize benchmarking practice as a two-stage evaluation: (i) computational interface geometry assessment (iRMSD/DockQ/RMSD plus confidence or PAE metrics) and (ii) experimental validation (binding, activation, and killing assays) [55].

Sequence- vs. structure-based generation for TCR-specific design. Sequence-only language models (e.g., ESM-2–style approaches) are highly effective for learning repertoire priors and generating human-like, developable TCR variants that respect statistical constraints of V/J usage, CDR composition, and overall fold compatibility [5]. However, because antigen specificity is determined by 3D interface geometry — including docking orientation and CDR-loop conformations — Sequence-only generation is typically insufficient *on its own* for reliably optimizing peptide–HLA recognition without an explicit structural validation layer. By contrast, structure-conditioned approaches (e.g., backbone/interface-conditioned sequence design or diffusion-style geometric generation) are better matched to antigen-specific TCR engineering because they can impose interface constraints directly (contact geometry, shape complementarity, steric feasibility) and localize flexibility to CDR loops while preserving framework stability [53]. In practice, the most effective TCR pipelines are hybrid: sequence models propose diverse, repertoire-conditioned candidates, while structure-aware models and/or structure predictors screen

and refine candidates for target-specific interface plausibility before experimental validation.

Biophysical quality is further evaluated through stability and energy-based metrics. Predicted folding free energy (ΔG), interface energy, and solvent accessibility are estimated using energy functions or molecular-dynamics–derived scoring [133]. Designs with unfavorable ΔG values, high steric clash scores, or unstable secondary-structure profiles are filtered out before experimental testing.

Binding performance is evaluated using both computational and experimental metrics. In silico docking and scoring functions estimate binding affinity, interface complementarity, and contact geometry between TCR and peptide–HLA complexes [93,128]. Experimentally, affinity and kinetics are measured using surface plasmon resonance, biolayer interferometry, or multimer staining, yielding dissociation constants (K_D), on-rates, and off-rates that quantify binding strength and stability [134,135].

Functional quality is assessed at the cellular level. Engineered T cells expressing generated receptors are evaluated by flow cytometry for surface expression and antigen-specific binding, by activation markers such as CD69 and CD25, and by cytokine secretion (e.g., IFN-γ, IL-2). Cytotoxic performance is quantified using chromium-release, luminescence-based killing assays, or live-cell imaging to measure selective lysis of antigen-positive targets [136,137].

Safety and specificity metrics are increasingly incorporated into evaluation frameworks. These include predicted or measured binding to panels of off-target peptides, cross-reactivity indices based on structural similarity, and functional assays against healthy-cell antigens [129,138,139]. Multi-objective scoring functions integrate affinity, specificity, stability, and safety into composite metrics that guide selection of candidates for experimental validation.

Together, these evaluation layers form a hierarchical metric system: sequence and structure plausibility, biophysical stability, binding strength, cellular function, and off-target risk. This multiscale benchmarking is essential for comparing generative models, validating design claims, and translating computational designs into clinically meaningful immune receptors.

Structure-aware networks further refine this capacity. Frameworks such as *ProteinMPNN-TCR*, *AlphaBind*, and *ImmuneDiffusion* explicitly encode geometric and energetic constraints, allowing the generation of receptors that maintain structural stability and realistic interface complementarity [32,87,88,140]. Diffusion-based models trained on crystallographic complexes of TCR–pMHC or antibody–antigen interactions learn the

conditional probability distribution of amino-acid arrangements within binding interfaces and can thus propose residues likely to enhance affinity without inducing steric conflict [141,142]. These approaches combine evolutionary information, structural priors, and physical constraints to produce receptor sequences that balance functional novelty with biophysical plausibility.

Although this section focuses on TCRs, the same generative design principles and evaluation frameworks extend naturally to antibody and nanobody engineering. Antibody and nanobody design have similarly benefited from generative architectures. Protein language models fine-tuned on antibody repertoires capture canonical framework and CDR motifs while permitting targeted diversification of paratopes [143,144]. Diffusion networks have been used to generate entire variable domains consistent with specific antigenic epitopes identified by cryo-electron microscopy or deep mutational scanning [145,146]. Generative sampling across latent manifolds defined by affinity, solubility, and expression metrics enables the creation of variant ensembles optimized for multiple objectives simultaneously. The resulting computationally derived antibodies and TCR mimetics extend the natural immune toolkit toward synthetic precision molecules with programmable binding properties [88,147].

5.2 Modular Optimization of Chimeric Antigen Receptors

Chimeric antigen receptors (CARs) are synthetic constructs that rewire immune recognition into an engineered signaling cascade. Each CAR consists of distinct functional modules; an extracellular binding domain, a hinge and transmembrane region, and one or more intracellular signaling motifs. The performance of a CAR depends on the integrated behavior of these modules, yet empirical optimization through domain swapping and screening is slow and labor-intensive [148]. Generative AI has introduced data-driven strategies capable of exploring this modular design space systematically [117].

Transformer-based models represent CAR components as compositional sequences that encode domain identity, positional order, and contextual interdependencies. By training on curated libraries of CAR constructs linked to phenotypic readouts, these architectures learn how variations in domain composition and arrangement modulate activation thresholds, cytokine signatures, and cellular persistence [149,150]. Conditional generation allows the creation of new CAR configurations optimized for desired functional signatures, such as reduced tonic signaling or enhanced metabolic fitness [151,152].

Reinforcement learning and active-learning algorithms further refine this design process. In these frameworks, model predictions are iteratively updated using experimental feedback from high-throughput CAR screening platforms, enabling convergence toward optimal

constructs [153,154,155]. Such feedback loops have already produced CARs with modified hinge lengths and co-stimulatory domain combinations that yield improved cytotoxic performance and diminished exhaustion markers, demonstrating how iterative design–build–test–learn cycles can directly optimize CAR function through adaptive computational–experimental feedback [156].

Generative modeling also supports the design of armored CARs, which incorporate payload modules such as cytokine secretion cassettes, chemokine receptors, or immune-checkpoint inhibitors. By embedding these additional modules within the same representational space, AI models can co-optimize receptor binding and paracrine modulation, resulting in constructs tailored for hostile tumor microenvironments [157,158]. Collectively, these developments illustrate how generative frameworks convert CAR engineering from heuristic assembly into a rational optimization problem solvable by machine learning.

5.3 Engineering Logic-Gated and Multiplexed Architectures

A further evolution of receptor design involves the encoding of logical operations into immune constructs. Logic-gated CARs and TCRs employ multi-antigen recognition to refine specificity and reduce off-target cytotoxicity [159]. Generative modeling enables systematic exploration of these multi-input architectures by representing antigens, linkers, and signaling modules within a shared latent space. In solid tumors, generative models are being used to design multispecific and logic-gated CAR architectures that improve selectivity and persistence while mitigating off-tumor toxicity [160]. Conditional diffusion or variational models can generate dual-specific binding domains whose cooperative interactions produce Boolean outcomes such as AND, OR, or NOT responses depending on antigen co-expression patterns [161].

These models facilitate computational optimization of interdomain spacing, linker composition, and binding affinity ratios required for balanced activation. By simulating dose–response landscapes across predicted antigen concentrations, AI systems can identify configurations that achieve strong tumor selectivity while sparing healthy tissues [2,162]. In addition, machine-learning-guided sampling of co-stimulatory domain combinations enables fine-tuning of intracellular signaling strength and timing [163,164]. Such multi-objective optimization integrates molecular recognition with system-level control, extending the reach of generative immunoengineering beyond molecular design to programmable cellular logic.

Multiplexed receptor architectures, which incorporate multiple signaling channels within a single cell, also benefit from generative approaches. Models trained on combinatorial

libraries of bispecific or tandem CARs learn statistical mappings between module composition and functional synergy [165,166]. The ability to generate thousands of candidate architectures *in silico* and evaluate them through predictive scoring significantly accelerates discovery, particularly for solid-tumor targets that require simultaneous recognition of multiple antigens.

5.4 Integrating Computational Design with Experimental Validation

The practical utility of generative receptor design depends on its integration with empirical validation. High-throughput display technologies, including yeast, phage, and mammalian systems, provide experimental evidence that grounds model predictions [167]. Single-cell transcriptomic and proteomic profiling captures downstream functional outcomes and supplies data for retraining generative models through active learning [168,169].

These advances have given rise to adaptive DBTL frameworks, in which computational generation and laboratory experimentation are connected through iterative feedback. Within such closed-loop systems, generative models propose receptor or construct candidates, automated biofoundries synthesize and screen them, and the resulting empirical data are reintegrated to update model priors. This adaptive coupling between in silico inference and in vitro validation forms the operational core of generative immunoengineering.

Closed-loop pipelines are emerging in which model-generated receptor or construct libraries are synthesized, expressed, and screened automatically. The resulting activity and expression data are re-entered into the model to update its priors, progressively improving generative accuracy [170]. These DBTL-driven feedback cycles enable continuous hypothesis generation, experimental execution, and model refinement to occur as a coordinated, self-correcting process under human oversight. This iterative design–build–test–learn cycle parallels the self-optimization processes characteristic of control theory and enables rapid convergence toward functional solutions.

Molecular-dynamics simulations and energy-based filtering further ensure that generated sequences respect physical constraints [171]. Structural relaxation, solvent accessibility analysis, and free-energy estimation help eliminate unstable or non-functional designs before synthesis [172]. When combined with automated DNA assembly and cell-based assays, these computational safeguards reduce experimental cost and improve hit rates. Collectively, these elements establish a hybrid experimental–computational ecosystem in which the DBTL cycle functions as the organizing principle linking algorithmic design to biological realization.

5.4.1 Computational Cost and Hardware Requirements

Generative immunoengineering pipelines vary widely in computational cost depending on model class, scale, and workflow stage. Sequence-based language models such as ESM-2 or immune-specific transformers can be fine-tuned and sampled on single high-memory GPUs or small multi-GPU servers, making them accessible to many academic and translational research labs [7]. Inference for sequence generation typically requires minutes to hours per batch, depending on model size and sampling strategy.

Structure-aware and diffusion-based models, including ProteinMPNN and RFdiffusion, impose higher computational demands [6,53]. Backbone-conditioned sequence design with ProteinMPNN is relatively lightweight and can be run on standard GPU workstations, whereas diffusion-based structure generation often requires multi-GPU systems and longer runtimes due to iterative denoising steps. Full 3D generative workflows combined with docking, energy minimization, and molecular-dynamics refinement can require access to high-performance computing clusters or cloud-based GPU infrastructure.

However, these costs are usually concentrated in the early design phase. Once candidate receptors or constructs are generated, downstream experimental screening dominates both time and expense. Many workflows therefore adopt a tiered strategy: lightweight sequence generation and filtering on local GPUs, followed by structure-based refinement and physics simulations on shared institutional clusters or commercial cloud platforms [104].

As a result, generative design is not restricted to supercomputing centers, but neither is it trivial for purely resource-limited clinical labs. In practice, most translational pipelines operate through collaborations with academic computing centers, biofoundries, or cloud providers, enabling scalable computation without requiring permanent in-house supercomputing infrastructure.

5.4.2 Hallucination Control and Physics-Based Filtering

Generative models can produce sequences or structures that appear statistically plausible yet violate physical or biochemical constraints, a phenomenon often referred to as “hallucination.” Importantly, hallucination does not refer only to low-confidence structural predictions (e.g., low pLDDT or high predicted aligned error). A more subtle and clinically relevant form involves high-confidence errors: sequences that neural networks predict with strong structural confidence, yet which fail under physical or thermodynamic constraints. These high-confidence hallucinations may satisfy learned statistical priors while exposing hydrophobic cores, promoting aggregation, destabilizing secondary structure, or forming energetically unfavorable interfaces under physiological conditions.

At the structural level, generated sequences are typically passed through structure-prediction or fold-validation models such as AlphaFold or AlphaFold-Multimer to assess folding confidence and interface plausibility using metrics such as pLDDT, predicted aligned error (PAE), and steric clash detection [5,173]. Candidates with low-confidence folds, backbone distortions, or unstable interface geometries are removed prior to downstream analysis.

Physics-based refinement further reduces hallucination risk. Many workflows apply energy minimization and structural relaxation using force-field–based tools such as Rosetta, OpenMM, or GROMACS [6,174,175]. These methods optimize side-chain packing, hydrogen bonding, and steric compatibility, eliminating candidates that collapse, unfold, or form high-energy conformations during relaxation. Docking simulations are often used to test whether generated receptors form stable complexes with their intended peptide–HLA targets and unstable or promiscuous complexes with off-target peptides [176].

Molecular-dynamics simulations provide an additional and particularly critical layer of validation for detecting high-confidence hallucinations. While neural-network confidence metrics (such as pLDDT) assess structural plausibility within learned statistical distributions, they do not directly evaluate thermodynamic stability in explicit solvent or dynamic physiological environments. MD relaxation, solvent-exposure analysis, and free-energy estimation can reveal buried hydrophobic residues that become exposed, unstable loop conformations, interface dissociation, or aggregation-prone surfaces that are not captured by static confidence scores [177,178]. In this sense, physics-based simulation serves as a necessary orthogonal filter, distinguishing designs that are statistically coherent from those that are thermodynamically viable.

Finally, experimental pipelines serve as the ultimate hallucination filter. Generated receptors are subjected to early-stage expression screening, surface localization assays, and binding tests [179,180,181]. Candidates that fail to fold, traffic, or bind appropriately are eliminated before functional assays. This layered strategy—statistical generation followed by structural prediction, physics-based refinement, and experimental filtering—reduces both low-confidence structural artifacts and high-confidence thermodynamic hallucinations, ensuring that generative models function not as unconstrained proposal engines but as components of a physically grounded design pipeline.

The convergence of computational and experimental pipelines also facilitates reproducibility and transparency. Standardized data formats, metadata capture, and open benchmarking of generative models are enabling comparative evaluation across laboratories [72,182]. These practices are essential for establishing trust in AI-generated

constructs and for supporting regulatory assessment of algorithmically designed therapeutics.

Overall, the integration of generative modeling, reinforcement optimization, and closed-loop DBTL validation defines a coherent framework for immune receptor and construct design. The transition from empirical mutagenesis to algorithmic generation compresses discovery timelines while expanding the accessible design space. As generative models increasingly integrate multimodal data linking molecular architecture to cellular outcomes, receptor design and phenotype programming are beginning to merge. This convergence marks the next stage of generative immunoengineering, where molecular construction and cellular behavior are optimized within a unified, adaptive learning system.

While receptor and construct design benefit from direct structural constraints and measurable binding outcomes, extending generative frameworks to cellular phenotypes introduces additional layers of complexity related to causality, regulatory feedback, and state stability.

6. Programming Cell Phenotypes with Generative Models

The capacity to engineer receptors and signaling modules has redefined the molecular architecture of immune cells, but the next frontier of generative design extends beyond receptor composition to the regulation of cellular phenotype [75,183]. Immune function is not determined solely by receptor specificity but by the emergent states of activation, metabolism, and gene regulation that arise within complex intracellular networks [1]. Generative models are now being adapted to learn and model these higher-order regulatory landscapes, providing a framework to propose candidate strategies for influencing differentiation, persistence, and functional polarization.

It is important to distinguish between generative modeling of cellular states and generative design of causal interventions. In contrast to receptor design, where structural constraints and biophysical validation allow relatively direct mapping between sequence and function, phenotype modeling operates in a higher-dimensional and more weakly supervised space. Most current generative phenotype frameworks learn statistical manifolds of transcriptional or epigenetic states derived primarily from observational single-cell datasets. These models can interpolate within learned state spaces and predict how a cell may resemble a desired phenotype, such as enhanced persistence or reduced exhaustion. However, resemblance does not imply causal sufficiency. The ability to reproduce the transcriptomic signature of a memory-like state does not guarantee that perturbing a predicted regulator will induce and stabilize that state in vivo. A causality gap therefore remains between state interpolation and

verified circuit-level reprogramming. Bridging this gap requires integration with perturb-seq datasets, mechanistic gene-network modeling, and closed-loop experimental validation.

6.1 Modeling the Cellular State Space

Immune cells exist within a high-dimensional state space defined by transcriptional, epigenetic, and metabolic variables that evolve dynamically in response to environmental stimuli. Traditional analytical frameworks, such as clustering or trajectory inference, describe these states retrospectively but do not predict how they can be reprogrammed [184,185]. Generative models such as variational autoencoders (VAEs), diffusion probabilistic models, and generative adversarial networks (GANs) provide a fundamentally different capability: they learn the underlying probability distribution of cellular states and can interpolate or sample from this learned manifold to predict unseen or engineered phenotypes [186,187].

While these models are powerful at learning associations between gene-expression patterns and phenotypic states, association alone is insufficient for true cellular programming. A transcriptional profile that resembles memory, persistence, or exhaustion does not imply that inducing that profile will *cause* a stable phenotypic transition. Programming immune cells ultimately requires intervening in causal regulatory networks, not merely reproducing correlational signatures. The central challenge therefore lies in moving from descriptive statements of the form *"this gene-expression state resembles memory"* to actionable causal hypotheses such as *"perturbing these regulators will induce and stabilize a memory phenotype."* This shift reframes generative phenotype modeling from a retrospective pattern-matching exercise into a forward-looking control problem centered on causal intervention.

A key limitation of current generative phenotype models is that the majority are trained predominantly on observational single-cell datasets, rather than systematic perturbation experiments. As a result, their ability to predict the outcomes of radical or out-of-distribution reprogramming remains largely unvalidated. Models may interpolate reliably within the manifold of observed states but fail when asked to extrapolate to phenotypes that require coordinated, multi-node regulatory intervention. Addressing this limitation will require deeper integration of perturb-seq data, causal inference frameworks, and closed-loop experimental validation to ensure that predicted interventions produce the intended phenotypic outcomes in living cells.

When trained on large-scale single-cell RNA sequencing (scRNA-seq) or ATAC-seq datasets, VAEs capture latent variables that correspond to biological processes such as activation, exhaustion, or memory differentiation [188,189,190]. These latent representations can be

manipulated to simulate transcriptional reprogramming trajectories. For instance, altering specific latent dimensions can emulate transitions from naïve to effector or from effector to exhausted states, revealing regulatory dependencies that govern these transitions. The learned latent manifold effectively approximates the probabilistic topology of the immune cell-state landscape, providing a computational analog to Waddington's epigenetic landscape but learned directly from data. Diffusion-based frameworks extend this capacity by modeling the stochastic evolution of gene-expression profiles, providing a generative account of cell-state dynamics over pseudo-temporal trajectories [191,192].

In parallel, multimodal models that integrate transcriptomic, proteomic, and metabolomic features are beginning to describe the coupled regulation of gene expression and metabolism in activated immune cells. Such models allow the generation of hypothetical phenotypes characterized by specific metabolic adaptations, cytokine secretion profiles, or migratory capacities [193,194]. By conditioning on environmental variables such as hypoxia, nutrient availability, or cytokine gradients, these frameworks can simulate how immune cells would adapt under diverse microenvironmental conditions [195]. However, the accuracy of such simulations remains constrained by the completeness and batch-corrected quality of training data, emphasizing the ongoing need for harmonized multimodal datasets.

Crucially, this conditioning-based generative strategy reflects a deeper shift in underlying assumptions. Rather than treating immune phenotypes as discrete, historically observed states, generative frameworks assume that immune behavior occupies a continuous functional landscape that can be explored and extrapolated beyond naturally sampled examples. This perspective enables the algorithmic exploration of rare, transient, or experimentally inaccessible immune states and reinforces the generative paradigm as one of design rather than discovery alone.

6.2 Generative Reprogramming and Perturbation Modeling

The translation of generative modeling from descriptive to prescriptive use involves connecting latent dimensions to actionable molecular interventions. Perturbation-based training strategies, such as those used in models like scGen and CPA (Compositional Perturbation Autoencoder), learn mappings between control and perturbed cellular states across thousands of experimental manipulations [196,197]. These models can then generate counterfactual predictions of how a given perturbation—such as gene knockout, cytokine exposure, or small-molecule treatment—would reprogram cellular transcriptional and proteomic profiles.

When coupled with CRISPR perturb-seq data, generative models can prioritize candidate sets of transcriptional regulators predicted to bias cellular trajectories toward desired phenotypic outcomes, subject to experimental validation. This framework has been applied to predict reprogramming strategies that induce T-cell memory phenotypes or reverse exhaustion-associated transcriptional signatures [198,199,200]. In macrophages and dendritic cells, similar models have been used to explore how combinations of signaling inputs reshape inflammatory versus tolerogenic polarization. These predictions can then guide targeted interventions using synthetic circuits, small molecules, or genome editing [201]. It is important to note that predictive fidelity depends on the coverage of the training manifold; models extrapolate reliably only within data-supported regions of perturbational space.

Integrating these models with reinforcement learning, further enables iterative optimization of intervention strategies. The algorithm explores a combinatorial action space of perturbations and uses feedback from simulated outcomes to propose the most effective intervention sequences. This approach reframes aspects of cellular reprogramming as an optimization problem, in which artificial intelligence can assist in identifying promising intervention sequences, allowing dynamic control of gene-regulatory networks rather than static modification of single targets.

6.3 Linking Generative Models to Synthetic Circuits

Generative modeling also provides a computational substrate for the design of synthetic gene circuits that enact desired cell-state transitions [202,203]. Once a target phenotype is defined in latent space such as resistance to exhaustion, enhanced persistence, or altered cytokine balance, AI models can identify candidate regulatory motifs or signaling pathways that need to be modulated to achieve that state. In this context, regulatory motifs may refer to either cis-regulatory DNA elements controlling transcriptional logic or dynamic feedback structures within signaling networks, depending on the level of abstraction. Synthetic biologists can then construct corresponding genetic circuits to implement these predicted control strategies.

For example, reinforcement learning coupled with gene-network simulations has been used to design circuit architectures that stabilize T-cell metabolic fitness by dynamically regulating glycolytic and oxidative pathways [204,205]. Diffusion-based generators trained on transcriptional responses to immune checkpoints have proposed feedback modules that mitigate activation-induced exhaustion [34,206]. Such designs translate the statistical regularities learned by generative models into actionable biological logic, closing the gap between abstract representation and physical implementation.

The integration of generative models with experimental libraries of promoters and enhancer elements further enables data-driven optimization of regulatory sequences. By learning mappings between sequence composition and expression amplitude or inducibility, generative models can propose synthetic regulatory elements that achieve precise transcriptional tuning within engineered immune cells [207]. This capability is especially valuable for balancing effector potency and safety in next-generation CAR-T or TCR-engineered therapies, where overactivation or premature exhaustion can compromise efficacy.

6.4 Toward Closed-Loop Phenotype Design

A defining feature of generative phenotype modeling is the potential for closed-loop optimization in which computational predictions are continuously refined through empirical feedback [208]. Integration with high-throughput perturbation platforms, time-lapse imaging, and multi-omic profiling allows real-time assessment of how engineered interventions reshape cellular states [209]. Data from each iteration is used to update generative priors, improving accuracy and adaptability.

This closed-loop paradigm parallels the design–build–test–learn cycles established for molecular engineering but operates at the systems level of cellular behavior [210,211]. Extending the molecular DBTL framework to the cellular systems level enables iterative refinement of both molecular components and emergent phenotypes within a unified feedback architecture. In such frameworks, the model functions as a control algorithm that continuously adjusts interventions to maintain desired phenotypic states. Implementing these adaptive DBTL frameworks requires standardized data formats and interoperable experimental protocols to ensure reproducibility and safe automation. These systems could eventually form the basis of autonomous adaptive immunoengineering platforms in which AI proposes genetic or pharmacological modifications, the laboratory executes them through automated microfluidic experimentation, and the resulting data retrain the model in real time.

Implementing such feedback architecture requires not only computational sophistication but also standardized experimental protocols and interoperable data formats. Advances in laboratory automation, robotic culture systems, and real-time single-cell monitoring are making these integrations feasible. As models become capable of predicting the dynamic responses of engineered immune cells, phenotype programming may evolve from a trial-and-error discipline into a continuous adaptive optimization process [212].

The use of generative models to program immune cell phenotypes represents a conceptual expansion of synthetic immunology. It extends the logic of receptor and construct design

into the realm of dynamic cellular behavior. By learning from the multidimensional data that describe activation, differentiation, and adaptation, AI systems can propose intervention strategies that achieve desired phenotypic equilibria with minimal experimental iteration [213]. The resulting convergence of generative modeling, perturbation analysis, and synthetic circuit design transforms cellular reprogramming from a descriptive science into a predictive and creative enterprise.

In this emerging framework, Immune cells may increasingly be treated as programmable entities whose behavior can be guided through data-informed and experimentally validated intervention strategies. This capacity to generate and stabilize beneficial phenotypes, whether through transcriptional modulation, metabolic rewiring, or synthetic gene networks, constitutes a central milestone on the path toward programmable immunity **(Figure 3)**.

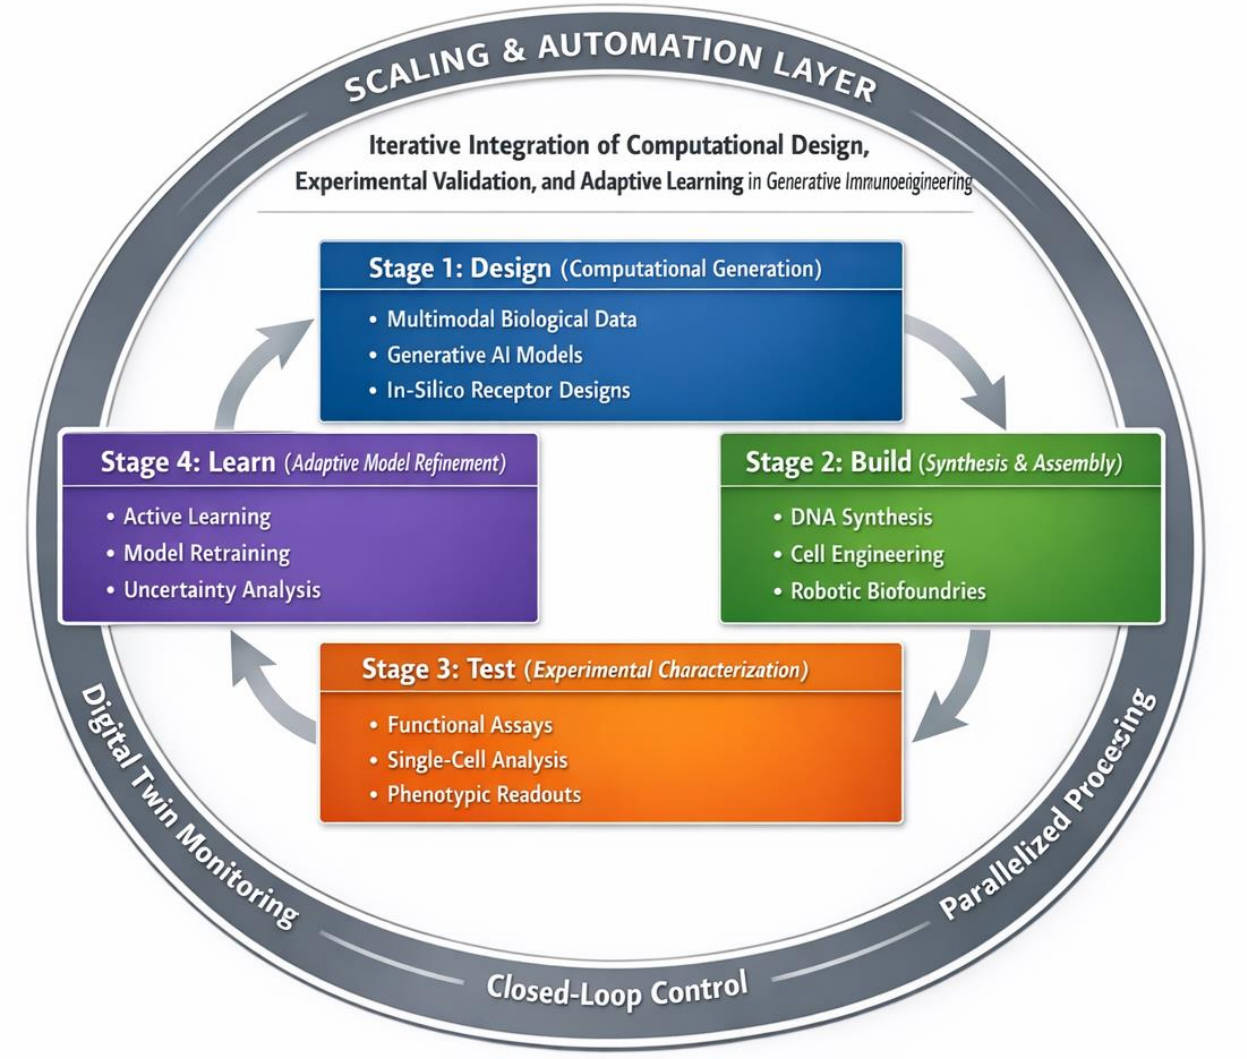

Figure 3. The closed-loop DBTL pipeline for generative immunoengineering.

7. The Design–Build–Test–Learn Loop at Scale

The maturation of generative immunoengineering depends not only on algorithmic sophistication but on the integration of computation, automation, and experimentation within a unified feedback architecture. The DBTL loop formalizes this integration as a recursive cycle in which hypotheses generated by artificial intelligence are iteratively realized, evaluated, and used to retrain the model [211,214]. At scale, DBTL transforms immunoengineering from a sequential workflow into a continuously adaptive system that fuses discovery, validation, and optimization into a single process.

7.1 The Design Phase: Generative Hypothesis Formation

In the generative paradigm, design constitutes a computational experiment in which the model explores the probability distribution of biological functions. Large foundation models trained on multi-omic, and structural datasets generate receptor sequences, circuit architectures, or cell-state perturbations that satisfy defined objective functions such as affinity, stability, specificity, metabolic resilience, and manufacturability [215,216].

Multi-objective reinforcement learning (MORL) and Pareto-front optimization are increasingly used to balance these criteria, ensuring that improvements in one property do not compromise another. For example, reinforcement agents can adjust generative sampling to favor constructs that maintain predicted folding stability while maximizing target binding and minimizing immunogenic epitopes [113,114,147]. Bayesian optimization frameworks quantify uncertainty across latent dimensions, guiding exploration toward regions of design space where the model's confidence is low but potential reward is high.

At the cellular level, generative models design intervention strategies that reprogram gene-regulatory networks or metabolic fluxes to achieve stable phenotypes. These designs may take the form of predicted transcription-factor combinations, circuit topologies, or epigenetic modifications [116,217]. In silico simulations using agent-based or ODE-based digital twins of immune cells allow evaluation of predicted designs before synthesis, effectively providing *pre-experimental validation* within the design phase itself [115]. This digital pre-screening step serves as an internal "virtual test phase," reducing wet-lab load while enhancing safety and design traceability within the DBTL framework.

7.2 The Build Phase: Automated Synthesis and Cellular Integration

The build phase converts digital designs into tangible biological constructs. Modern biofoundries employ modular, high-throughput synthesis pipelines that integrate robotic liquid handling, automated cloning, and barcoded sample tracking. DNA assembly methods such as Golden Gate, Gibson, and enzymatic ligation-independent cloning enable parallel production of thousands of constructs in standardized vectors [218,219].

In immunoengineering, the build step involves integrating these synthetic constructs into cellular systems. CRISPR/Cas, base-editing, and transposon-mediated delivery methods allow targeted insertion of designed sequences into immune-cell genomes, often at safe-harbor loci that permit consistent expression [220]. Microfluidic electroporation and viral-vector platforms have been optimized for parallel processing of primary T or NK cells, enabling libraries of engineered variants to be generated under controlled conditions [221]. Each design instance is annotated with its origin, parameters, and vector architecture,

allowing traceable linkage between computational proposal and biological realization. This traceability is critical for both reproducibility and regulatory compliance, ensuring transparent lineage from digital design to physical construct.

Emerging cell-free systems provide an intermediate validation layer. DNA templates or mRNA constructs can be expressed in vitro to assay folding, binding, or signaling activity before introduction into living cells [222,223]. These rapid screening layers reduce the cost and biosafety burden of testing AI-generated sequences. Integration with laboratory-information management systems (LIMS) ensures that metadata, sequence provenance, and performance metrics flow seamlessly back into the digital design environment [224].

7.3 The Test Phase: High-Dimensional and Multiscale Evaluation

Testing represents the sensory layer of the DBTL system, translating experimental outcomes into quantitative metrics for model refinement. High-throughput display systems—yeast, phage, or mammalian—provide initial binding and expression readouts for receptor libraries [103]. Flow cytometry and surface plasmon resonance quantify affinity and kinetic constants, while single-cell assays capture functional endpoints such as cytokine release, proliferation, or exhaustion markers [225].

Recent developments in multi-omic screening have expanded the granularity of testing. Single-cell RNA sequencing, proteomic barcoding, and metabolomic profiling characterize thousands of engineered cells simultaneously, revealing how synthetic constructs reshape global cellular states [226]. Spatial transcriptomics and live-cell imaging provide contextual information about cell–cell interactions, trafficking, and synapse formation.

These high-dimensional data streams are analyzed through unsupervised embedding and graph-based clustering to extract latent features representing functional archetypes. Statistical coupling between design parameters and phenotypic readouts allows causal inference about which molecular features drive performance. Importantly, uncertainty quantification metrics guide the selection of candidates for deeper mechanistic analysis, ensuring that the test phase not only validates designs but also enriches the informational value of subsequent learning cycles [227]. Incorporating Bayesian calibration and explainable modeling frameworks can further ensure that performance gains are mechanistically interpretable rather than purely correlational.

At scale, the DBTL loop reframes the laboratory itself as a sensory interface for artificial intelligence. The Test phase is not merely a validation checkpoint but the primary data-generation mechanism that determines how rapidly and accurately models can learn. The throughput, quality, and dimensionality of experimental measurements therefore directly

bound the learning rate of the entire generative system. As a consequence, advances in generative immunoengineering are intrinsically coupled to advances in high-content experimental automation, including single-cell multi-omics, live-cell imaging, and parallelized functional assays. In this regime, the principal bottleneck may shift from computational design capacity to empirical characterization, positioning experimental infrastructure as a rate-limiting component of algorithmic progress [228].

7.4 The Learn Phase: Model Updating and Active Reinforcement

The learn phase closes the feedback loop by converting empirical results into updated model parameters. Instead of static retraining, modern systems employ *online learning* architectures in which models ingest experimental data in near real time [154,229]. Each new data batch adjusts the model's latent embeddings, probability weights, and uncertainty estimates, progressively aligning computational predictions with biological reality.

Active learning strategies determine which experiments would most efficiently reduce model uncertainty. The algorithm selects a subset of candidates predicted to yield the highest information gain, focusing experimental resources on the most informative regions of design space [9]. Reinforcement learning further couples the model to the physical system: successful experimental outcomes increase reward signals that bias subsequent generative sampling toward productive directions.

In practice, these adaptive feedback create a cyber-physical learning organism. An integrated system in which computational and biological components co-evolve. Each iteration not only refines the model's internal representation but also generates new empirical priors that expand its capacity to generalize [230]. The result is exponential acceleration of discovery efficiency, with each loop yielding designs of higher predicted performance and lower variance between simulation and experiment. This iterative refinement embodies a data-driven analogue of biological evolution (variation, selection, and retention) executed within an AI-governed experimental ecosystem.

7.5 Automation, Data Infrastructure, and Self-Optimizing Biofoundries

At scale, DBTL becomes inseparable from automation. Modern biofoundries combine robotics, microfluidics, and advanced data orchestration to enable continuous closed-loop experimentation. AI design servers communicate directly with robotic assembly lines through standardized APIs, initiating synthesis and testing sequences without manual intervention [231,232]. Real-time sensor data—temperature, reagent usage, cell viability, expression levels—are streamed to cloud infrastructures that synchronize model updates.

Digital twins of the laboratory simulate physical processes in silico, allowing predictive scheduling, error correction, and adaptive re-prioritization of experimental tasks. These twins maintain a live correspondence between virtual and real experiments, permitting instantaneous recalibration when deviations occur [233]. Integration with edge-computing modules enables local decision-making, reducing latency between data acquisition and design refinement.

Data interoperability is central to scaling. Standardized ontologies (e.g., SBOL, AnIML, MIFlowCyt) and metadata schemas ensure that information from diverse instruments and facilities can be aggregated for cross-institutional learning. Cloud-native data lakes equipped with version control and provenance tracking store raw and processed datasets, supporting reproducibility and regulatory auditing [234,235]. Standardization of both data semantics and experiment-level metadata remains a bottleneck, underscoring the importance of community-driven interoperability initiatives.

As these infrastructures mature, self-optimizing biofoundries are emerging facilities where generative AI orchestrates the entire pipeline from molecular design to functional evaluation. Such systems can autonomously evolve improved receptor variants, optimize circuit architectures, and fine-tune culture parameters to enhance yield or stability [236]. Over time, the accumulated data acts as an institutional memory, enabling transfer learning across projects and the continuous improvement of both algorithms and experimental protocols.

7.6 Integrative and Translational Implications

The large-scale implementation of DBTL frameworks in immunoengineering signifies a structural reorganization of how biological knowledge is produced. Instead of discrete projects defined by static hypotheses, research becomes a dynamic optimization process governed by real-time feedback [237,238]. Generative models no longer operate as isolated analytical tools but as components of an evolving experimental ecosystem.

At the translational level, scalable DBTL systems accelerate the path from computational concept to clinical candidate. By systematically linking receptor sequence, cell phenotype, and manufacturing parameters, these platforms can identify predictive markers of efficacy and safety early in development [239,240]. Closed-loop optimization also supports adaptive manufacturing, where process parameters are adjusted algorithmically to maintain product quality in response to real-time analytics.

Ultimately, the convergence of generative AI, automation, and scalable experimentation transform immunoengineering into a continuously learning infrastructure. Each iteration

expands the collective intelligence encoded in both digital models and biological systems, gradually approaching a regime where the boundaries between *designing immunity*, *testing immunity*, and *learning immunity* become indistinguishable. This architecture represents the operational foundation of programmable immunity, translating theoretical possibility into a self-refining experimental reality. In this architecture, the DBTL framework functions not merely as an engineering workflow but as a new epistemology for biological design—one where knowledge generation, model evolution, and therapeutic innovation proceed as an inseparable continuum.

8. Translational Opportunities and Clinical Outlook

Generative immunoengineering represents a fundamental reorganization of how cell-based therapies are conceived, tested, and manufactured. What was once a linear sequence—spanning discovery, optimization, and production is becoming an adaptive continuum in which computation, experimentation, and clinical translation are tightly coupled through iterative feedback [76]. Within this architecture, the immune system is no longer treated solely as a biological entity to be modulated but as a programmable substrate whose molecular and cellular functions can be designed, validated, and continuously refined through artificial intelligence. This conceptual shift recasts translational medicine as a dynamic learning process, one in which biology and computation evolve in synchrony. This paradigm reframes therapeutic development as a bidirectional learning system in which clinical data, molecular modeling, and experimental outcomes continuously inform one another, effectively creating a feedback-coupled translational pipeline.

8.1 Versioning, Algorithmic Dossiers, and the Identity of Adaptive Cell Therapies

As generative immunoengineering systems mature, translational challenges shift from single-entity approval toward governance of adaptive therapeutic platforms. In this context, versioning emerges as a critical regulatory concept. When a generative model is retrained and produces an updated construct—for example, a “CAR-T version 2.0”—the central question becomes whether this update represents an incremental refinement within a validated design envelope or the creation of a fundamentally new biological entity.

Addressing this distinction requires moving beyond static molecular descriptions toward algorithmic dossiers that track design lineage, training data provenance, conditioning variables, and objective functions across iterations. Rather than treating each generated construct as an isolated product, the therapeutic identity can be framed at the platform level, with predefined boundaries specifying permissible variation in sequence, signaling architecture, or phenotypic output. Updates that remain within these validated boundaries

may be considered iterative improvements, whereas changes that alter mechanism of action, state-space occupancy, or risk profile would trigger additional preclinical or clinical evaluation.

This challenge is amplified by the biological reality that immune cells exist within a high-dimensional state space defined by transcriptional, epigenetic, and metabolic variables that evolve dynamically in response to environmental stimuli [241]. Small algorithmic or construct-level modifications can therefore propagate nonlinearly across cellular states, underscoring the need for version-aware oversight that integrates computational change logs with empirical phenotypic validation. Framing adaptive immune therapies through the lens of versioned platforms, rather than static entities, provides a principled pathway for balancing continuous learning with regulatory rigor.

Despite rapid progress in generative modeling and experimental validation, the single largest barrier preventing AI-designed TCRs from entering clinical trials in 2025 is not algorithmic performance but regulatory-grade validation of safety and specificity. While models can now generate high-affinity, structurally plausible receptors, current pipelines lack standardized, regulator-accepted frameworks for demonstrating absence of dangerous cross-reactivity at the scale required for first-in-human testing. Predicting rare but catastrophic off-target recognition remains extremely difficult due to incomplete coverage of the human peptidome, limited negative datasets, and imperfect structural generalization [124,242,243]. As a result, translation is constrained less by design capability than by the need for scalable, trusted preclinical safety evaluation systems that combine large off-target libraries, structure-aware screening, and high-throughput functional assays under regulatory oversight.

At the foundation of this transformation is the capacity of generative models to accelerate discovery and optimization with unprecedented efficiency. Large-scale protein and cellular language models trained on structural, sequence, and binding-affinity data can propose millions of receptor or circuit variants that satisfy pre-defined biophysical and functional constraints. Multi-objective optimization—combining Bayesian inference, reinforcement learning, and evolutionary search—allows competing design criteria such as affinity, folding stability, and manufacturability to be reconciled within a single probabilistic framework [147,216]. When coupled to high-throughput synthesis and functional screening, these algorithms transform receptor design from an empirical search into a statistically guided exploration of sequence space. Early analyses suggest that such workflows can reduce the experimental burden by an order of magnitude, compressing timelines that once spanned years into a matter of weeks, though such acceleration remains contingent on access to high-quality multimodal training data and harmonized experimental standards, while

preserving, and in some cases improving success rates in identifying viable therapeutic constructs.

Beyond speed, the generative paradigm introduces the possibility of genuine personalization. Patient-derived molecular data including tumor transcriptomes, HLA genotypes, immune-repertoire sequencing, and single-cell multi-omics can serve as conditioning variables for model inference [213]. This enables the generation of individualized T-cell receptors or chimeric antigen receptors that are predicted to engage patient-specific neoantigens while minimizing self-reactivity. In principle, these digital blueprints can be synthesized directly into autologous lymphocytes or natural-killer cells within closed, automated manufacturing systems [244]. Crucially, the same framework permits *adaptive therapy*, continuous molecular monitoring, through circulating tumor DNA, antigen-escape profiling, or cytokine dynamics can be fed back to retraining models and update constructs in response to disease evolution. This continuous feedback loop effectively transforms treatment into a dynamic control process, aligning therapeutic pressure with tumor or immune-escape kinetics in near real time. Therapy thus becomes a moving equilibrium, a co-adaptive process in which the treatment learns from the patient as much as the patient responds to the treatment [245].

Realizing this adaptive vision requires new regulatory and clinical-trial architecture. Current frameworks presuppose fixed molecular entities, yet generative therapeutics are inherently dynamic. Regulatory bodies have begun developing guidance under the principles of Good Machine Learning Practice, emphasizing transparency, explainability, and post-deployment surveillance. Adaptive or platform trials may replace static designs, allowing algorithmically derived construct revisions within predefined boundaries. Version-controlled documentation of model parameters, training data, and validation outcomes will constitute an "algorithmic dossier," analogous to the chemistry-manufacturing-controls documentation required for biologics [246,247]. Post-market oversight is expected to include continuous monitoring for model drift and periodic re-certification of retrained algorithms. Together, these mechanisms will form the basis of a new discipline, regulatory bioinformatics, dedicated to the governance of learning systems in medicine. Such governance will likely integrate algorithmic explainability metrics, model-card disclosures, and standardized digital audit trails to ensure accountability throughout the therapeutic life cycle.

Translation from digital design to clinical-grade production further depends on automation and digital infrastructure. AI-integrated biofoundries now link computational design servers directly to robotic assembly, viral-vector packaging, and closed-system cell expansion, creating an unbroken digital thread from algorithmic proposal to physical manufacture. Each

construct carries a persistent identifier linking its computational origin to its production batch, ensuring traceability across the product life cycle [248]. Digital-twin bioreactors simulate nutrient gradients, cytokine signaling, and metabolic flux in real time, adjusting culture conditions through reinforcement-learning controllers to preserve cell viability and phenotypic stability. Multi-omic sensors monitoring transcriptomic, impedance, and metabolic signatures feed continuous data into these control layers, allowing predictive correction of deviations before product quality is compromised. Such cyber-physical feedback transforms Good Manufacturing Practice environments from static production lines into adaptive learning systems that improve with every run [249,250]. Collectively, these infrastructures redefine Good Manufacturing Practice (GMP) as a dynamic rather than static framework, one in which quality is maintained through continuous sensing, prediction, and correction rather than retrospective testing.

Although oncology remains the initial testing ground, the principles of generative immunoengineering are applicable across diverse therapeutic landscapes. In solid tumors, generative models are being used to design multispecific and logic-gated CAR architectures that improve selectivity and persistence while mitigating off-tumor toxicity. In autoimmunity, the same computational logic can be applied to the design of regulatory-T-cell or dendritic-cell circuits that restore immune tolerance without systemic suppression. Regenerative medicine may benefit from macrophage or tolerogenic antigen-presenting-cell designs optimized for cytokine balance and metabolic resilience, promoting tissue repair and graft acceptance **(Supplementary Table S3)** [67,117, 251-259]. Similar approaches extend to infectious-disease preparedness, where rapid, model-driven updates to antigen or receptor design could allow immune interventions to evolve as quickly as the pathogens they target.

Ensuring the safety and interpretability of algorithmically derived constructs remains paramount. Multi-layered validation pipelines combine *in silico* prediction, explainable-AI analysis, and empirical verification. Generative outputs are screened for immunogenic motifs, structural instability, and potential off-target binding before synthesis. These in-silico safeguards are complemented by human-in-the-loop review protocols to prevent over-reliance on automated predictions in high-risk therapeutic contexts. Attention-based visualization and feature-attribution mapping identify sequence regions most influential in the model's predictions, while uncertainty quantification provides calibrated confidence estimates. Empirical assays ranging from multiplex peptide arrays to single-cell cytotoxicity screens serve as orthogonal tests of computational accuracy [260,261]. To formalize accountability, proposed standards for Algorithmic Documentation Files will record the architecture, data provenance, and performance metrics of each deployed model, enabling reproducibility and auditability across regulatory jurisdictions [262].

Economic and infrastructural considerations are equally transformative. The high up-front computational investment in data curation and model training is counterbalanced by substantial downstream savings from reduced screening, accelerated iteration, and automated manufacturing. Distributed biofoundries connected through secure cloud infrastructures could enable regional or hospital-based production of autologous therapies, reducing logistic complexity and cold-chain dependency. Achieving this vision will require harmonized digital quality-management systems and interoperable GMP documentation across sites. Ensuring equitable access to these infrastructures, particularly in low- and middle-income regions will be critical to prevent a widening translational divide in personalized immunotherapy. Health-economic evaluation frameworks must evolve to recognize the amortized value of continuously improving algorithms and the outcomes they enable, transitioning from static cost-effectiveness assessments toward performance-linked reimbursement models [263].

The ethical, legal, and social implications of generative immunoengineering must evolve alongside its technical capabilities. The use of clinical and genomic data for model training necessitates explicit consent frameworks specifying secondary and longitudinal data use. Questions of intellectual property whether ownership resides in datasets, model architectures, or generated sequences require global policy coordination. Dual-use risks are real; systems capable of designing potent therapeutic receptors could, in principle, be misused to generate immune-evasive or pathogenic molecules. Safeguarding measures analogous to existing biosecurity treaties will be essential, as will equitable access to computational infrastructure to prevent concentration of capability within a few technologically privileged centers. Open repositories, distributed compute alliances, and internationally governed consortia can help ensure that the benefits of programmable immunity are shared globally rather than confined to specific regions or sectors [264,265].

Taken together, these translational trajectories delineate the emergence of a new therapeutic paradigm. In the near term, AI-optimized CAR and TCR constructs with integrated safety and quality-monitoring frameworks are poised to enter early-phase trials. Over the next decade, standardized digital biofoundries and adaptive regulatory pipelines are likely to support the extension of this technology to autoimmune, transplant, and regenerative contexts [117]. Beyond that horizon lies the prospect of a continuously learning therapeutic ecosystem in which each patient outcome refines the generative models that guide subsequent design. In this future, therapeutic innovation and biological understanding become inseparable, and the immune system itself is reimagined as a programmable interface between computation and life. The defining hallmark of AI-enabled generative design of immune cells and receptors for programmable immunity [266]. This synthesis of

adaptive intelligence and living matter marks not only a technological milestone but a conceptual redefinition of medicine itself, one in which therapy, learning, and evolution converge within a single generative continuum.

9. Ethical, Regulatory, and Societal Implications of Generative Immunoengineering

The capacity to design immune cells and receptors through generative artificial intelligence represents both a profound scientific advance and a consequential ethical turning point. The very features that make this technology transformative; its speed, adaptability, and autonomy, also challenge the traditional mechanisms through which biomedical innovation has been governed. As generative models begin to influence the structure of experimental inquiry, the criteria of clinical validation, and the distribution of therapeutic access, they introduce a new layer of moral and regulatory responsibility. This creates an obligation to design not only biological systems but also the systems of oversight that will ensure their safe and equitable use [267,268]. This dual responsibility to engineer biology and the ethics that govern it defines an emerging field of generative bioethics evolving in parallel with generative biology itself.

9.1 Embedded Ethics in Generative Immunoengineering

Ethical governance in generative immunoengineering cannot be treated as an external review layer applied after technical development. Instead, ethics must be embedded directly within the DBTL loop, co-evolving with algorithms, data, and engineered cells. In programmable immunity, governance itself must be programmable, adaptive, and technically instantiated within the same workflows that generate biological function.

At the design stage, ethical constraints can be formalized as optimization objectives rather than qualitative guidelines. Fairness and bias metrics may be incorporated directly into model loss functions or reward signals, ensuring that generated receptors or cell programs do not systematically privilege specific genetic backgrounds, tissue contexts, or demographic groups. Similarly, algorithmic explainability can be treated as a required design output, with interpretable representations and causal attributions generated alongside candidate constructs rather than as post hoc audits.

At the data and infrastructure level, embedded ethics motivates the adoption of privacy-preserving architectures as default components of generative pipelines. Federated learning, secure multiparty computation, and differential privacy enable models to learn from distributed clinical and genomic datasets without centralizing sensitive patient information.

Integrating these approaches into biofoundry-linked DBTL systems ensures that scale and learning efficiency do not come at the expense of privacy, consent, or data sovereignty.

During testing and learning, ethical governance continues through continuous monitoring of safety, uncertainty, and downstream risk. Experimental feedback not only refines performance predictions but also updates ethical risk assessments, allowing models to adaptively restrict exploration of designs associated with elevated immunogenicity, off-target activity, or misuse potential. In this framework, ethics is not a static compliance checkpoint but a dynamic control layer that constrains and guides generative exploration throughout the lifecycle of adaptive immune therapies.

The following paragraphs examine how this embedded ethical logic manifests across data stewardship, interpretability, regulatory oversight, and societal impact.

At the ethical level, the most immediate concern involves the use of patient-derived data for model training and algorithmic conditioning. The effectiveness of generative immunoengineering depends on access to large, high-quality datasets encompassing genomic, proteomic, and clinical information. Yet the aggregation of such data often derived from identifiable biological samples raises questions about consent, ownership, and longitudinal use. Current consent models, designed for discrete studies, are poorly suited to the continuous learning paradigm of AI-driven research. Dynamic or “evergreen” consent frameworks, allowing participants to renew, modify, or revoke data permissions as models evolve, may become essential to align data use with individual autonomy [269,270,271]. In clinical settings, such adaptive consent must be coupled with ongoing feedback between patients and therapeutic models, ensuring that participants maintain agency as their biological and algorithmic profiles co-evolve. Likewise, new institutional mechanisms will be needed to recognize participants not merely as data donors but as contributors to the generative process, with potential claims to benefit-sharing or acknowledgment. Including under-represented populations in training datasets will also be critical to mitigate demographic bias and ensure global generalizability of AI-driven immune design.

Transparency and explainability constitute a second ethical axis. The interpretability of generative outputs is crucial for both scientific trust and clinical safety. While recent advances in attention mapping, saliency analysis, and uncertainty quantification have improved insight into how models generate biological designs, the epistemic gap between statistical pattern recognition and causal biological reasoning persists. Regulators, clinicians, and researchers must therefore treat algorithmic explainability not as an optional feature but as a moral imperative. Documentation of model lineage, training data provenance, and version history should be considered integral to ethical disclosure,

comparable in importance to reporting methods in experimental science [272,273]. Bridging this interpretive gap will require hybrid frameworks that combine mechanistic modeling with data-driven generation so that algorithmic creativity remains biologically intelligible. Without this transparency, the reproducibility crisis that has affected other fields could extend into synthetic immunology, undermining confidence in AI-mediated design.

Regulatory institutions now face the challenge of overseeing entities that are not static products but evolving systems. Conventional approval pathways, designed for fixed molecular entities, cannot easily accommodate therapeutic platforms that learn from new data and autonomously propose novel constructs [274]. Emerging frameworks under the rubric of *Good Machine Learning Practice* (GMLP) attempt to address this tension by introducing algorithmic change-control mechanisms, documentation standards, and real-time performance monitoring **(Figure 4)**. In the context of generative immunoengineering, such regulation will likely require integration of *algorithmic dossiers* into the chemistry-manufacturing-controls infrastructure of GMP [275]. Each generative model may be regarded as a "living protocol" subject to continuous validation and regulatory auditing. This convergence of computational oversight and biomanufacturing control represents a defining shift in the governance of biomedical AI and will require cross-trained professionals fluent in both algorithmic governance and GMP compliance.

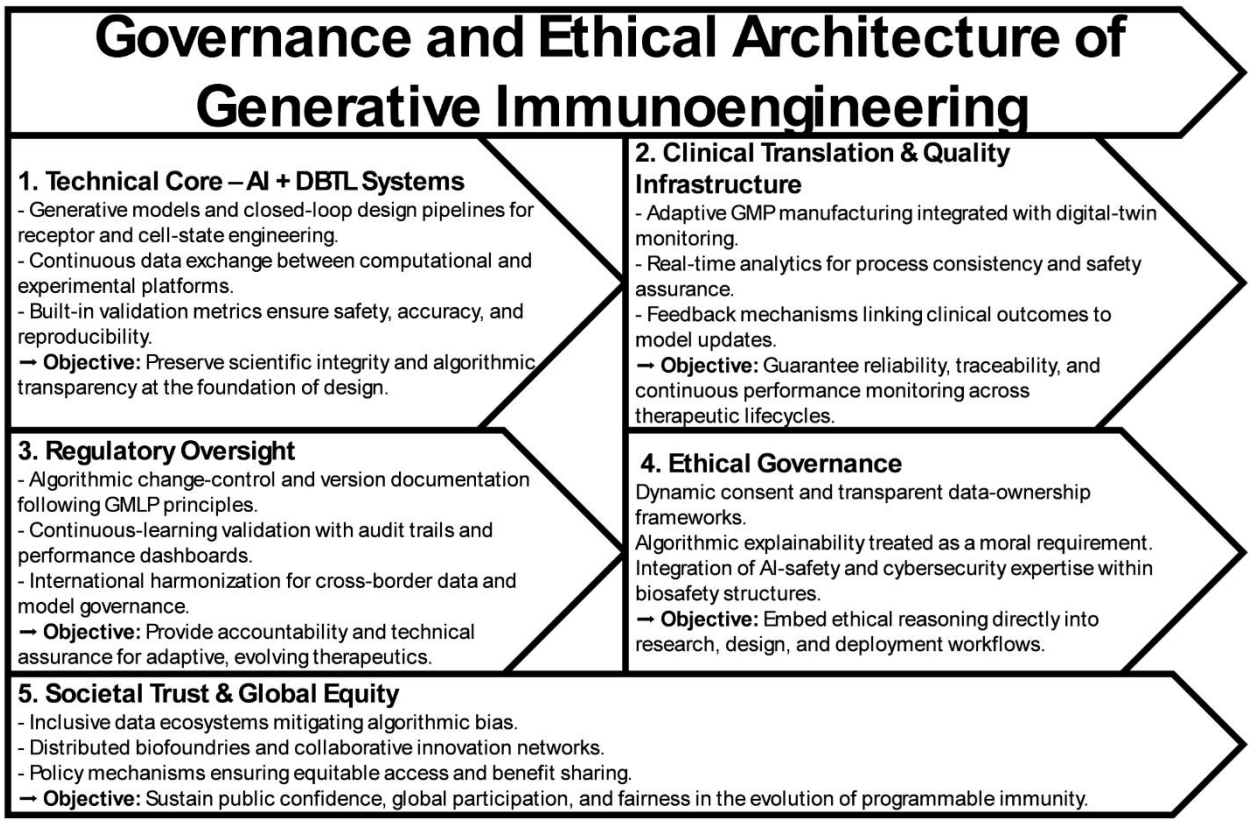

Figure 4. Governance Framework for Generative Immunoengineering From Technical Foundations to Societal Oversight

A further dimension of ethical responsibility arises from dual-use potential and biosecurity. The same generative architectures that optimize immune recognition could, in principle, be repurposed to design immune-evasive pathogens, synthetic toxins, or receptor-binding antagonists. As algorithmic tools diffuse through open scientific ecosystems, the line between beneficial and hazardous application becomes porous. International biosecurity regimes, historically focused on material agents and physical laboratories, must therefore expand to encompass informational biosecurity—the governance of digital models, codebases, and data pipelines capable of generating biological functions. Building multi-layered safeguards that combine technical containment, federated architectures, and ethical licensing will be central to future risk management. Developing secure access frameworks such as controlled model release, tiered licensing, and federated training will be essential to balance open scientific exchange with risk mitigation [268,276]. Ethical oversight committees within research institutions should incorporate expertise in AI safety and cybersecurity alongside traditional biosafety representation.

The societal implications of generative immunoengineering extend beyond bioethics into the political economy of biomedical innovation. Because generative design relies on large

computational infrastructure and proprietary datasets, there is a risk that technological capacity and therefore therapeutic opportunity will become concentrated within a small number of institutions and nations. Without deliberate intervention, the “AI divide” in healthcare could mirror and amplify existing inequities in access to genomic medicine. Counteracting this trend will require international coordination of data-sharing standards, open access repositories for pre-trained biological models, and collaborative licensing arrangements that enable low-resource regions to participate in the development and deployment of AI-driven therapeutics [277,278]. Equitable access is not merely a matter of distributive justice but of scientific robustness. Diversity in training data enhances generalizability and reduces model bias, thereby improving global therapeutic safety. Moreover, algorithmic asymmetry in data ownership and computational access risks consolidating economic power among a few institutions, creating new forms of biomedical dependency that demand policy-level correction.

The epistemological implications are equally significant. Generative immunoengineering reconfigures the relationship between hypothesis and experiment, transforming discovery into an iterative dialogue between algorithmic inference and empirical validation [279,280]. This shift blurs the historical boundary between knowledge generation and technological fabrication, forcing a reconsideration of what counts as “understanding” in biology. If a model can design a receptor that functions optimally without the designer fully comprehending the underlying causal grammar, the locus of scientific agency moves from the human investigator to the hybrid system of human and machine. In this configuration, agency becomes distributed, a co-production of human intention and algorithmic inference, raising new questions about authorship, accountability, and epistemic responsibility. Such transformations invite reflection on the nature of explanation, accountability, and authorship in the age of algorithmic biology. The challenge for future scientific culture will be to ensure that interpretability and conceptual insight evolve alongside performance and automation [279].

Finally, the integration of generative immunoengineering into clinical and societal systems will demand new forms of governance that are anticipatory rather than reactive. Ethical frameworks should be embedded from the outset of model development, not retrofitted in response to controversy. Cross-disciplinary oversight bringing together immunologists, data scientists, ethicists, clinicians, and patient representatives should guide decisions about data use, design objectives, and therapeutic deployment. International consortia may serve as coordinating bodies to establish shared principles for algorithmic transparency, data equity, and biosafety. As generative biology becomes increasingly autonomous, the human responsibility for defining its boundaries and purposes becomes more, not less, essential.

The future of programmable immunity will thus depend not only on the sophistication of its algorithms but also on the moral architecture of the institutions that steward them [268]. Only through such ethically adaptive governance can programmable immunity mature into a discipline that safeguards both biological integrity and societal trust in the era of generative medicine.

Conclusion

The convergence of generative artificial intelligence, cellular engineering, and immunology marks a transformative moment in the life sciences. What began as an attempt to enhance receptor design has matured into a new framework for understanding and shaping biology itself. Through generative modeling, immune repertoires, signaling networks, and cellular phenotypes can now be explored as dynamic design spaces rather than static natural entities. This shift dissolves the traditional boundaries between discovery and fabrication, between observing life and programming it. It redefines biology as an editable, self-informing system, one in which learning becomes a property of both the organism and the methods that study it.

AI-enabled generative immunoengineering creates a continuum that links molecular design, phenotypic programming, and automated biomanufacturing within integrated feedback systems. The design–build–test–learn cycle converts cell therapy development from a linear experimental sequence into a continuously adaptive process. Each iteration strengthens the intelligence of the system by transforming empirical outcomes into computational insight, allowing therapeutic design and biological understanding to co-evolve. In this model, medicine becomes a learning enterprise, guided by algorithms that refine their predictions through every patient and experiment. Such integration heralds the emergence of a “living laboratory” paradigm, where computation, experimentation, and clinical feedback operate as one self-optimizing network.

Translationally, this paradigm redefines both the structure and the pace of biomedical innovation. The ability to generate, validate, and deploy patient-specific immune receptors within compressed timeframes makes precision immunotherapy responsive to individual and population-level dynamics. Automated manufacturing environments equipped with digital twins and real-time analytics promise production systems that not only reproduce validated protocols but improve upon them through continuous optimization. As regulatory frameworks evolve to accommodate algorithmic validation and adaptive approval, the distinction between discovery, manufacturing, and clinical deployment will diminish, giving rise to a self-improving therapeutic ecosystem embedded within healthcare itself. In the long view, this same architecture may extend beyond immunology, linking generative genomics,

regenerative medicine, and neural interface design into a unified science of programmable biology.

At the same time, this technological acceleration intensifies questions of ethics, governance, and equity. The capacity to design immunity at will requires oversight systems that are as adaptive as the technologies they regulate. Consent, data ownership, intellectual property, and algorithmic transparency must be treated as integral components of the research architecture rather than external constraints. The future of generative immunoengineering will depend on maintaining equilibrium between innovation and accountability, openness and security, personalization and fairness. The sophistication of our moral and institutional design must keep pace with the sophistication of our computational tools. Ethical foresight must therefore evolve as dynamically as the algorithms themselves, ensuring that creativity and responsibility remain inseparable.

Ultimately, generative immunoengineering invites a new way of thinking about intervention in biology. It envisions a medicine in which intelligence, human and artificial, acts as a creative partner in the shaping of immune function. The concept of programmable immunity captures this synthesis. A vision of therapeutic science that is predictive, personalized, and continuously learning, yet guided by ethical foresight and social responsibility. It signals the beginning of a new epistemology in which the design of life and the design of knowledge become one continuous act. The challenge ahead is not only to design better immune cells but to design wiser systems; scientific, ethical, and societal through which such power can be directed toward the enduring goal of human well-being.

List of abbreviations:

AI – Artificial Intelligence
AND – Logical AND gate
ATAC-seq – Assay for Transposase-Accessible Chromatin using sequencing
BCR – B-Cell Receptor
BI – Biolayer Interferometry
CAR – Chimeric Antigen Receptor
CAR-T – Chimeric Antigen Receptor T cell
CD – Cluster of Differentiation
CDR – Complementarity-Determining Region
CPA – Compositional Perturbation Autoencoder
CRISPR – Clustered Regularly Interspaced Short Palindromic Repeats
DBTL – Design–Build–Test–Learn
DeepImmuno – Deep-learning-based immunogenicity prediction framework

DNA – Deoxyribonucleic Acid
DockQ – Docking Quality score
ESM – Evolutionary Scale Modeling
GAN – Generative Adversarial Network
GNN – Graph Neural Network
GPU – Graphics Processing Unit
HLA – Human Leukocyte Antigen
IEDB – Immune Epitope Database
iRMSD – Interface Root-Mean-Square Deviation
K_D – Dissociation Constant
MD – Molecular Dynamics
MHC – Major Histocompatibility Complex
MORL – Multi-Objective Reinforcement Learning
MSA – Multiple Sequence Alignment
NetMHCpan – Neural-network-based MHC binding prediction tool
NOT – Logical NOT gate
ODE – Ordinary Differential Equation
OAS – Observed Antibody Space
PAE – Predicted Aligned Error
pLDDT – Predicted Local Distance Difference Test
pMHC – Peptide–Major Histocompatibility Complex
ProtT5 – Protein T5 language model
RFdiffusion – RoseTTAFold diffusion-based protein generation framework
RMSD – Root-Mean-Square Deviation
RL – Reinforcement Learning
scRNA-seq – Single-Cell RNA Sequencing
SPR – Surface Plasmon Resonance
TCR – T-Cell Receptor
TCR-T – T-Cell Receptor-Engineered T cell
VAE – Variational Autoencoder
VDJdb – V(D)J Database

Author Contributions:

Conceptualization: O.A., M.N.; Methodology: M.N.; Formal analysis: O.A., M.N.; Investigation: M.N.; Resources: M.N.; Writing—original draft preparation: O.A., M.N.; Writing—review and editing: O.A., M.N.; Visualization: M.N.; Supervision: M.N.; Project

administration: O.A., M.N. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflicts of interest.

Acknowledgment: None

AI-Declaration:

Google Gemini was used for limited language editing in a small number of sections. All scientific content, interpretations, and conclusions were developed solely by the authors. All figures are original, were prepared by M.N. for this manuscript, with the assistance of AI-based visualization tools.

Supplementary Material

Supplementary material associated with this article has been published online and is available at: Link to the DOI

References

1. Navarro Quiroz R, Villarreal Camacho J, Zarate Penata E, et al. (2025) Multiscale information processing in the immune system. *Front Immunol* 16: 1563992.
2. Dewaker V, Morya VK, Kim YH, et al. (2025) Revolutionizing oncology: the role of Artificial Intelligence (AI) as an antibody design, and optimization tools. *Biomark Res* 13: 52.
3. Rafelski SM, Theriot JA (2024) Establishing a conceptual framework for holistic cell states and state transitions. *Cell* 187: 2633-2651.
4. Leary AY, Scott D, Gupta NT, et al. (2024) Designing meaningful continuous representations of T cell receptor sequences with deep generative models. *Nat Commun* 15: 4271.
5. Jumper J, Evans R, Pritzel A, et al. (2021) Highly accurate protein structure prediction with AlphaFold. *Nature* 596: 583-589.
6. Dauparas J, Anishchenko I, Bennett N, et al. (2022) Robust deep learning-based protein sequence design using ProteinMPNN. *Science* 378: 49-56.
7. Rives A, Meier J, Sercu T, et al. (2021) Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proc Natl Acad Sci U S A* 118.
8. Ibrahim M, Khalil YA, Amirrajab S, et al. (2025) Generative AI for synthetic data across multiple medical modalities: A systematic review of recent developments and challenges. *Comput Biol Med* 189: 109834.

9. Zhang P, Wei L, Li J, et al. (2024) Artificial intelligence-guided strategies for next-generation biological sequence design. *Natl Sci Rev* 11: nwae343.
10. Ertelt M, Moretti R, Meiler J, et al. (2025) Self-supervised machine learning methods for protein design improve sampling but not the identification of high-fitness variants. *Sci Adv* 11: eadr7338.
11. Krapp LF, Meireles FA, Abriata LA, et al. (2024) Context-aware geometric deep learning for protein sequence design. *Nat Commun* 15: 6273.
12. Ma JL, H.; Hu, Y.; Huang, J. (2025) Physicochemically Informed Dual-Conditioned Generative Model of T-Cell Receptor Variable Regions for Cellular Therapy. *arxiv*.
13. Chouleur T, Etchegaray C, Villain L, et al. (2025) A strategy for multimodal integration of transcriptomics, proteomics, and radiomics data for the prediction of recurrence in patients with IDH-mutant gliomas. *Int J Cancer* 157: 573-587.
14. Berson E, Chung P, Espinosa C, et al. (2024) Unlocking human immune system complexity through AI. *Nat Methods* 21: 1400-1402.
15. Gao S, Fang A, Huang Y, et al. (2024) Empowering biomedical discovery with AI agents. *Cell* 187: 6125-6151.
16. O'Donnell TJ, Kanduri C, Isacchini G, et al. (2024) Reading the repertoire: Progress in adaptive immune receptor analysis using machine learning. *Cell Syst* 15: 1168-1189.
17. Liu YT, Zhang LL, Jiang ZY, et al. (2025) Applications of Artificial Intelligence in Biotech Drug Discovery and Product Development. *MedComm (2020)* 6: e70317.
18. Katoh H, Komura D, Furuya G, et al. (2023) Immune repertoire profiling for disease pathobiology. *Pathol Int* 73: 1-11.
19. Ou Y, Guo S (2023) Safety risks and ethical governance of biomedical applications of synthetic biology. *Front Bioeng Biotechnol* 11: 1292029.
20. Kong Y, Li J, Zhao X, et al. (2024) CAR-T cell therapy: developments, challenges and expanded applications from cancer to autoimmunity. *Front Immunol* 15: 1519671.
21. de Greef PC, Oakes T, Gerritsen B, et al. (2020) The naive T-cell receptor repertoire has an extremely broad distribution of clone sizes. *Elife* 9.
22. Tsahouridis O, Xu M, Song F, et al. (2025) The landscape of CAR-engineered innate immune cells for cancer immunotherapy. *Nat Cancer* 6: 1145-1156.
23. Gangwal A, Ansari A, Ahmad I, et al. (2024) Generative artificial intelligence in drug discovery: basic framework, recent advances, challenges, and opportunities. *Front Pharmacol* 15: 1331062.
24. Johnson JAI, Bergman DR, Rocha HL, et al. (2025) Human interpretable grammar encodes multicellular systems biology models to democratize virtual cell laboratories. *Cell* 188: 4711-4733 e4737.
25. Weber CR, Rubio T, Wang L, et al. (2022) Reference-based comparison of adaptive immune receptor repertoires. *Cell Rep Methods* 2: 100269.
26. Brown N, Fiscato M, Segler MHS, et al. (2019) GuacaMol: Benchmarking Models for de Novo Molecular Design. *J Chem Inf Model* 59: 1096-1108.
27. Sanchez-Lengeling B, Aspuru-Guzik A (2018) Inverse molecular design using machine learning: Generative models for matter engineering. *Science* 361: 360-365.

28. Bjerregaard A, Groth PM, Hauberg S, et al. (2025) Foundation models of protein sequences: A brief overview. *Curr Opin Struct Biol* 91: 103004.
29. Fox DR, Taveneau C, Clement J, et al. (2025) Code to complex: AI-driven de novo binder design. *Structure* 33: 1631-1642.
30. Wang M, Patsenker J, Li H, et al. (2025) Supervised fine-tuning of pre-trained antibody language models improves antigen specificity prediction. *PLoS Comput Biol* 21: e1012153.
31. Chen LT, Quinn Z, Dumas M, et al. (2025) Target sequence-conditioned design of peptide binders using masked language modeling. *Nat Biotechnol*.
32. Gasser HC, Rajan A, Alfaro JA (2025) A novel decoding strategy for ProteinMPNN to design with less visibility to cytotoxic T-lymphocytes. *Comput Struct Biotechnol J* 27: 3693-3703.
33. Mei A, Letscher KP, Reddy S (2024) Engineering next-generation chimeric antigen receptor-T cells: recent breakthroughs and remaining challenges in design and screening of novel chimeric antigen receptor variants. *Curr Opin Biotechnol* 90: 103223.
34. Zhang J, Che Y, Liu R, et al. (2025) Deep learning-driven multi-omics analysis: enhancing cancer diagnostics and therapeutics. *Brief Bioinform* 26.
35. Sandberg TE, Salazar MJ, Weng LL, et al. (2019) The emergence of adaptive laboratory evolution as an efficient tool for biological discovery and industrial biotechnology. *Metab Eng* 56: 1-16.
36. Zhu L, Mou W, Hong C, et al. (2024) The Evaluation of Generative AI Should Include Repetition to Assess Stability. *JMIR Mhealth Uhealth* 12: e57978.
37. Mirakhori F, Niazi SK (2025) Harnessing the AI/ML in Drug and Biological Products Discovery and Development: The Regulatory Perspective. *Pharmaceuticals (Basel)* 18.
38. Teng F, Cui T, Zhou L, et al. (2024) Programmable synthetic receptors: the next-generation of cell and gene therapies. *Signal Transduct Target Ther* 9: 7.
39. Weissenow K, Rost B (2025) Are protein language models the new universal key? *Curr Opin Struct Biol* 91: 102997.
40. Chen Y, Wang Z, Zeng X, et al. (2023) Molecular language models: RNNs or transformer? *Brief Funct Genomics* 22: 392-400.
41. Brandes N, Ofer D, Peleg Y, et al. (2022) ProteinBERT: a universal deep-learning model of protein sequence and function. *Bioinformatics* 38: 2102-2110.
42. Alley EC, Khimulya G, Biswas S, et al. (2019) Unified rational protein engineering with sequence-based deep representation learning. *Nat Methods* 16: 1315-1322.
43. Lupo U, Sgarbossa D, Bitbol AF (2022) Protein language models trained on multiple sequence alignments learn phylogenetic relationships. *Nat Commun* 13: 6298.
44. Lin Z, Akin H, Rao R, et al. (2023) Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* 379: 1123-1130.
45. Chen Y, Xu Y, Liu D, et al. (2024) An end-to-end framework for the prediction of protein structure and fitness from single sequence. *Nat Commun* 15: 7400.

46. Wang M, Patsenker J, Li H, et al. (2024) Language model-based B cell receptor sequence embeddings can effectively encode receptor specificity. *Nucleic Acids Res* 52: 548-557.
47. Ornes S (2022) Researchers turn to deep learning to decode protein structures. *Proc Natl Acad Sci U S A* 119: e2202107119.
48. Heil BJ, Hoffman MM, Markowetz F, et al. (2021) Reproducibility standards for machine learning in the life sciences. *Nat Methods* 18: 1132-1135.
49. Justyna M, Zirbel C, Antczak M, et al. (2025) Graph neural network and diffusion model for modeling RNA interatomic interactions. *Bioinformatics* 41.
50. Soleymani F, Paquet E, Viktor HL, et al. (2024) Structure-based protein and small molecule generation using EGNN and diffusion models: A comprehensive review. *Comput Struct Biotechnol J* 23: 2779-2797.
51. Baek M, DiMaio F, Anishchenko I, et al. (2021) Accurate prediction of protein structures and interactions using a three-track neural network. *Science* 373: 871-876.
52. Wu KE, Yang KK, van den Berg R, et al. (2024) Protein structure generation via folding diffusion. *Nat Commun* 15: 1059.
53. Watson JL, Juergens D, Bennett NR, et al. (2023) De novo design of protein structure and function with RFdiffusion. *Nature* 620: 1089-1100.
54. Fernandez-Quintero ML, Pomarici ND, Loeffler JR, et al. (2020) T-Cell Receptor CDR3 Loop Conformations in Solution Shift the Relative Valpha-Vbeta Domain Distributions. *Front Immunol* 11: 1440.
55. Quast NP, Abanades B, Guloglu B, et al. (2025) T-cell receptor structures and predictive models reveal comparable alpha and beta chain structural diversity despite differing genetic complexity. *Commun Biol* 8: 362.
56. Bennett NR, Coventry B, Goreshnik I, et al. (2023) Improving de novo protein binder design with deep learning. *Nat Commun* 14: 2625.
57. Cao L, Coventry B, Goreshnik I, et al. (2022) Design of protein-binding proteins from the target structure alone. *Nature* 605: 551-560.
58. Ingraham JB, Baranov M, Costello Z, et al. (2023) Illuminating protein space with a programmable generative model. *Nature* 623: 1070-1078.
59. Ribeiro-Filho HV, Jara GE, Guerra JVS, et al. (2024) Exploring the potential of structure-based deep learning approaches for T cell receptor design. *PLoS Comput Biol* 20: e1012489.
60. Meehl MM, Immadisetty K, Trivedi VD, et al. (2025) Computational structural optimization enhances IL13Ralpha2 - B7-H3 tandem CAR T cells to overcome antigen-heterogeneity-mediated tumor escape. *Mol Ther* 33: 4968-4987.
61. Qiu S, Chen J, Wu T, et al. (2024) CAR-Toner: an AI-driven approach for CAR tonic signaling prediction and optimization. *Cell Res* 34: 386-388.
62. Nunes FVM, Behrens LMP, Weimer RD, et al. (2025) Deep learning methods and applications in single-cell multimodal data integration. *Mol Omics*.

63. Liu 刘俊伟 J, Cen 岑萧萍 X, Yi 伊晨昕 C, et al. (2025) Challenges in AI-driven Biomedical Multimodal Data Fusion and Analysis. *Genomics Proteomics Bioinformatics* 23.
64. Wu X, Yang X, Dai Y, et al. (2024) Single-cell sequencing to multi-omics: technologies and applications. *Biomark Res* 12: 110.
65. Lamiable A, Champetier T, Leonardi F, et al. (2023) Revealing invisible cell phenotypes with conditional generative modeling. *Nat Commun* 14: 6386.
66. Agarwal AA, Harrang J, Noble D, et al. (2025) AlphaBind, a domain-specific model to predict and optimize antibody-antigen binding affinity. *MAbs* 17: 2534626.
67. Karthikeyan D, Bennett SN, Reynolds AG, et al. (2025) Conditional generation of real antigen-specific T cell receptor sequences. *Nat Mach Intell* 7: 1494-1509.
68. Lawton ML, Inge MM, Blum BC, et al. (2024) Multiomic profiling of chronically activated CD4+ T cells identifies drivers of exhaustion and metabolic reprogramming. *PLoS Biol* 22: e3002943.
69. Iu DS, Maya J, Vu LT, et al. (2024) Transcriptional reprogramming primes CD8+ T cells toward exhaustion in Myalgic encephalomyelitis/chronic fatigue syndrome. *Proc Natl Acad Sci U S A* 121: e2415119121.
70. Hudson D, Fernandes RA, Basham M, et al. (2023) Can we predict T cell specificity with digital biology and machine learning? *Nat Rev Immunol* 23: 511-521.
71. Leggieri PA, Liu Y, Hayes M, et al. (2021) Integrating Systems and Synthetic Biology to Understand and Engineer Microbiomes. *Annu Rev Biomed Eng* 23: 169-201.
72. Leipzig J, Nust D, Hoyt CT, et al. (2021) The role of metadata in reproducible computational research. *Patterns (N Y)* 2: 100322.
73. Tule S, Foley G, Boden M (2024) Do protein language models learn phylogeny? *Brief Bioinform* 26.
74. Vieira MC, Palm AE, Stamper CT, et al. (2023) Germline-encoded specificities and the predictability of the B cell response. *PLoS Pathog* 19: e1011603.
75. Javdan SB, Deans TL (2021) Design and development of engineered receptors for cell and tissue engineering. *Curr Opin Syst Biol* 28.
76. Wong WW, Lim WA (2023) Golden age of immunoengineering. *Immunol Rev* 320: 4-9.
77. Rees AR (2020) Understanding the human antibody repertoire. *MAbs* 12: 1729683.
78. Villanueva-Flores F, Sanchez-Villamil JI, Garcia-Atutxa I (2025) Publisher Correction: AI-driven epitope prediction: a systematic review, comparative analysis, and practical guide for vaccine development. *NPJ Vaccines* 10: 209.
79. Elfatimi E, Lekbach Y, Prakash S, et al. (2025) Artificial intelligence and machine learning in the development of vaccines and immunotherapeutics-yesterday, today, and tomorrow. *Front Artif Intell* 8: 1620572.
80. Shugay M, Bagaev DV, Zvyagin IV, et al. (2018) VDJdb: a curated database of T-cell receptor sequences with known antigen specificity. *Nucleic Acids Res* 46: D419-D427.
81. Zhang W, Wang L, Liu K, et al. (2020) PIRD: Pan Immune Repertoire Database. *Bioinformatics* 36: 897-903.

82. Textor J, Buytenhuijs F, Rogers D, et al. (2023) Machine learning analysis of the T cell receptor repertoire identifies sequence features of self-reactivity. *Cell Syst* 14: 1059-1073 e1055.
83. Greenshields-Watson A, Abanades B, Deane CM (2024) Investigating the ability of deep learning-based structure prediction to extrapolate and/or enrich the set of antibody CDR canonical forms. *Front Immunol* 15: 1352703.
84. Ostrovsky-Berman M, Frankel B, Polak P, et al. (2021) Immune2vec: Embedding B/T Cell Receptor Sequences in R (N) Using Natural Language Processing. *Front Immunol* 12: 680687.
85. Sidhom JW, Larman HB, Pardoll DM, et al. (2021) DeepTCR is a deep learning framework for revealing sequence concepts within T-cell repertoires. *Nat Commun* 12: 1605.
86. Li L, Peng X, Batliwala M, et al. (2023) Crystal structures of MHC class I complexes reveal the elusive intermediate conformations explored during peptide editing. *Nat Commun* 14: 5020.
87. Sumida KH, Nunez-Franco R, Kalvet I, et al. (2024) Improving Protein Expression, Stability, and Function with ProteinMPNN. *J Am Chem Soc* 146: 2054-2061.
88. Meng F, Zhou N, Hu G, et al. (2024) A comprehensive overview of recent advances in generative models for antibodies. *Comput Struct Biotechnol J* 23: 2648-2660.
89. Lin V, Cheung M, Gowthaman R, et al. (2025) TCR3d 2.0: expanding the T cell receptor structure database with new structures, tools and interactions. *Nucleic Acids Res* 53: D604-D608.
90. Schmirler R, Heinzinger M, Rost B (2024) Fine-tuning protein language models boosts predictions across diverse tasks. *Nat Commun* 15: 7407.
91. Drost F, An Y, Bonafonte-Pardas I, et al. (2024) Multi-modal generative modeling for joint analysis of single-cell T cell receptor and gene expression data. *Nat Commun* 15: 5577.
92. Abramson J, Adler J, Dunger J, et al. (2024) Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature* 630: 493-500.
93. Deleuran SN, Nielsen M (2025) NetTCR-struc, a structure driven approach for prediction of TCR-pMHC interactions. *Front Immunol* 16: 1616328.
94. Lourenco A, Subramanian A, Spencer R, et al. (2025) Protein CREATE enables closed-loop design of de novo synthetic protein binders. *bioRxiv*.
95. Yang J, Lal RG, Bowden JC, et al. (2025) Active learning-assisted directed evolution. *Nat Commun* 16: 714.
96. Gao Y, Gao Y, Wu S, et al. (2025) Weakly supervised peptide-TCR binding prediction facilitates neoantigen identification. *Cell Syst* 16: 101403.
97. Pertseva M, Follonier O, Scarcella D, et al. (2024) TCR clustering by contrastive learning on antigen specificity. *Brief Bioinform* 25.
98. Saadat M, Zare-Mirakabad F, Masoudi-Nejad A, et al. (2024) HLAPepBinder: An Ensemble Model for The Prediction Of HLA-Peptide Binding. *Iran J Biotechnol* 22: e3927.

99. Gioia D, Bertazzo M, Recanatini M, et al. (2017) Dynamic Docking: A Paradigm Shift in Computational Drug Discovery. *Molecules* 22.
100. Guner Yilmaz OZ, Doruker P, Kurkcuoglu O (2025) A Computationally Efficient Method to Generate Plausible Conformers for Ensemble Docking and Binding Free Energy Calculations. *J Chem Inf Model* 65: 8137-8157.
101. Bielska W, Jaszczyszyn I, Dudzic P, et al. (2025) Applying computational protein design to therapeutic antibody discovery - current state and perspectives. *Front Immunol* 16: 1571371.
102. Sverchkov Y, Craven M (2017) A review of active learning approaches to experimental design for uncovering biological networks. *PLoS Comput Biol* 13: e1005466.
103. Slavny P, Hegde M, Doerner A, et al. (2024) Advancements in mammalian display technology for therapeutic antibody development and beyond: current landscape, challenges, and future prospects. *Front Immunol* 15: 1469329.
104. Ching T, Himmelstein DS, Beaulieu-Jones BK, et al. (2018) Opportunities and obstacles for deep learning in biology and medicine. *J R Soc Interface* 15.
105. Wossnig L, Furtmann N, Buchanan A, et al. (2024) Best practices for machine learning in antibody discovery and development. *Drug Discov Today* 29: 104025.
106. Rajagopal N, Choudhary U, Tsang K, et al. (2024) Deep learning-based design and experimental validation of a medicine-like human antibody library. *Brief Bioinform* 26.
107. Dixon T, MacPherson D, Mostofian B, et al. (2022) Predicting the structural basis of targeted protein degradation by integrating molecular dynamics simulations with structural mass spectrometry. *Nat Commun* 13: 5884.
108. Hati S, Bhattacharyya S (2016) Incorporating modeling and simulations in undergraduate biophysical chemistry course to promote understanding of structure-dynamics-function relationships in proteins. *Biochem Mol Biol Educ* 44: 140-159.
109. Hamamsy T, Barot M, Morton JT, et al. (2023) Learning sequence, structure, and function representations of proteins with language models. *bioRxiv*.
110. Dibaeinia P, Ojha A, Sinha S (2025) Interpretable AI for inference of causal molecular relationships from omics data. *Sci Adv* 11: eadk0837.
111. Davidsen K, Olson BJ, DeWitt WS, 3rd, et al. (2019) Deep generative models for T cell receptor protein sequences. *Elife* 8.
112. Isacchini G, Sethna Z, Elhanati Y, et al. (2020) Generative models of T-cell receptor sequences. *Phys Rev E* 101: 062414.
113. Al-Jumaily A, Mukaidaisi M, Vu A, et al. (2023) Examining multi-objective deep reinforcement learning frameworks for molecular design. *Biosystems* 232: 104989.
114. Wang J, Zhu F (2024) Multi-objective molecular generation via clustered Pareto-based reinforcement learning. *Neural Netw* 179: 106596.
115. Niarakis A, Laubenbacher R, An G, et al. (2024) Immune digital twins for complex human pathologies: applications, limitations, and challenges. *NPJ Syst Biol Appl* 10: 141.
116. Zrimec J, Fu X, Muhammad AS, et al. (2022) Controlling gene expression with deep generative design of regulatory DNA. *Nat Commun* 13: 5099.

117. Shahzadi M, Rafique H, Waheed A, et al. (2024) Artificial intelligence for chimeric antigen receptor-based therapies: a comprehensive review of current applications and future perspectives. *Ther Adv Vaccines Immunother* 12: 25151355241305856.
118. Croce G, Lani R, Tardivon D, et al. (2025) Phage display enables machine learning discovery of cancer antigen-specific TCRs. *Sci Adv* 11: eads5589.
119. Fast E, Dhar M, Chen B (2023) TAPIR: a T-cell receptor language model for predicting rare and novel targets. *bioRxiv*.
120. Semilietof A, Stefanidis E, Gray-Gaillard E, et al. (2024) Preclinical model for evaluating human TCRs against chimeric syngeneic tumors. *J Immunother Cancer* 12.
121. Liu B, Greenwood NF, Bonzanini JE, et al. (2025) Design of high-specificity binders for peptide-MHC-I complexes. *Science* 389: 386-391.
122. Johansen KH, Wolff DS, Scapolo B, et al. (2025) De novo-designed pMHC binders facilitate T cell-mediated cytotoxicity toward cancer cells. *Science* 389: 380-385.
123. Linette GP, Stadtmauer EA, Maus MV, et al. (2013) Cardiovascular toxicity and titin cross-reactivity of affinity-enhanced T cells in myeloma and melanoma. *Blood* 122: 863-871.
124. Cameron BJ, Gerry AB, Dukes J, et al. (2013) Identification of a Titin-derived HLA-A1-presented peptide as a cross-reactive target for engineered MAGE A3-directed T cells. *Sci Transl Med* 5: 197ra103.
125. Foldvari Z, Knetter C, Yang W, et al. (2023) A systematic safety pipeline for selection of T-cell receptors to enter clinical use. *NPJ Vaccines* 8: 126.
126. Reynisson B, Alvarez B, Paul S, et al. (2020) NetMHCpan-4.1 and NetMHCIIpan-4.0: improved predictions of MHC antigen presentation by concurrent motif deconvolution and integration of MS MHC eluted ligand data. *Nucleic Acids Res* 48: W449-W454.
127. Li G, Iyer B, Prasath VBS, et al. (2021) DeepImmuno: deep learning-empowered prediction and generation of immunogenic peptides for T-cell immunity. *Brief Bioinform* 22.
128. Bradley P (2023) Structure-based prediction of T cell receptor:peptide-MHC interactions. *Elife* 12.
129. Ishii K, Davies JS, Sinkoe AL, et al. (2023) Multi-tiered approach to detect autoimmune cross-reactivity of therapeutic T cell receptors. *Sci Adv* 9: eadg9845.
130. Rollins ZA, Faller R, George SC (2022) Using molecular dynamics simulations to interrogate T cell receptor non-equilibrium kinetics. *Comput Struct Biotechnol J* 20: 2124-2133.
131. Mazhibiyeva A, Pham TT, Pats K, et al. (2025) Bridging prediction and reality: Comprehensive analysis of experimental and AlphaFold 2 full-length nuclear receptor structures. *Comput Struct Biotechnol J* 27: 1998-2013.
132. Yin R, Ribeiro-Filho HV, Lin V, et al. (2023) TCRmodel2: high-resolution modeling of T cell receptor recognition using deep learning. *Nucleic Acids Res* 51: W569-W576.
133. Krishna R, Wang J, Ahern W, et al. (2024) Generalized biomolecular modeling and design with RoseTTAFold All-Atom. *Science* 384: eadl2528.

134. Robertson IB, Mulvaney R, Dieckmann N, et al. (2024) Tuning the potency and selectivity of ImmTAC molecules by affinity modulation. *Clin Exp Immunol* 215: 105-119.
135. Sinclair A, Kramer S, Reinhart C, et al. (2026) Beyond sequence similarity: ML-powered identification of pHLA off-targets for TCR-mimic antibodies using high throughput binding kinetics. *MAbs* 18: 2601360.
136. Wellach K, Riemer AB (2025) Highly sensitive live-cell imaging-based cytotoxicity assay enables functional validation of rare epitope-specific CTLs. *Front Immunol* 16: 1558620.
137. Foulke JG, Chen L, Chang H, et al. (2024) Optimizing Ex Vivo CAR-T Cell-Mediated Cytotoxicity Assay through Multimodality Imaging. *Cancers (Basel)* 16.
138. Banerjee A, Pattinson DJ, Wincek CL, et al. (2025) T cell receptor cross-reactivity prediction improved by a comprehensive mutational scan database. *Cell Syst* 16: 101345.
139. Sharma G, Teng F, Round J, et al. (2025) Comprehensive self-antigen screening to assess cross-reactivity in promiscuous T-cell receptors. *Front Immunol* 16: 1719827.
140. Tang X, Dai H, Knight E, et al. (2024) A survey of generative AI for de novo drug design: new frontiers in molecule and protein generation. *Brief Bioinform* 25.
141. Borrman T, Pierce BG, Vreven T, et al. (2021) High-throughput modeling and scoring of TCR-pMHC complexes to predict cross-reactive peptides. *Bioinformatics* 36: 5377-5385.
142. Seo SY, Rhee JK (2025) TCR-epiDiff: solving dual challenges of TCR generation and binding prediction. *Bioinformatics* 41: i125-i132.
143. Kalemati M, Noroozi A, Shahbakhsh A, et al. (2024) ParaAntiProt provides paratope prediction using antibody and protein language models. *Sci Rep* 14: 29141.
144. Ghanbarpour A, Jiang M, Foster D, et al. (2023) Structure-free antibody paratope similarity prediction for in silico epitope binning via protein language models. *iScience* 26: 106036.
145. Pruvost T, Mathieu M, Dubois S, et al. (2023) Deciphering cross-species reactivity of LAMP-1 antibodies using deep mutational epitope mapping and AlphaFold. *MAbs* 15: 2175311.
146. Mason DM, Reddy ST (2024) Predicting adaptive immune receptor specificities by machine learning is a data generation problem. *Cell Syst* 15: 1190-1197.
147. Abeer A, Urban NM, Weil MR, et al. (2024) Multi-objective latent space optimization of generative molecular design models. *Patterns (N Y)* 5: 101042.
148. Lindner SE, Johnson SM, Brown CE, et al. (2020) Chimeric antigen receptor signaling: Functional consequences and design implications. *Sci Adv* 6: eaaz3223.
149. Guedan S, Calderon H, Posey AD, Jr., et al. (2019) Engineering and Design of Chimeric Antigen Receptors. *Mol Ther Methods Clin Dev* 12: 145-156.
150. Castellanos-Rueda R, Wang KK, Forster JL, et al. (2025) Dissecting the role of CAR signaling architectures on T cell activation and persistence using pooled screens and single-cell sequencing. *Sci Adv* 11: eadp4008.

151. Smirnov S, Mateikovich P, Samochernykh K, et al. (2024) Recent advances on CAR-T signaling pave the way for prolonged persistence and new modalities in clinic. *Front Immunol* 15: 1335424.
152. Alsaieedi AA, Zaher KA (2025) Tracing the development of CAR-T cell design: from concept to next-generation platforms. *Front Immunol* 16: 1615212.
153. Sutanto HF, D. (2025) Integrating artificial intelligence into small molecule development for precision cancer immunomodulation therapy. *npj Drug Discovery* 2.
154. Du Q, Wang H, Jiang B, et al. (2025) Advancing genetic engineering with active learning: theory, implementations and potential opportunities. *Brief Bioinform* 26.
155. Ferdous S, Shihab IF, Chowdhury R, et al. (2024) Reinforcement learning-guided control strategies for CAR T-cell activation and expansion. *Biotechnol Bioeng* 121: 2868-2880.
156. Chandrasekar V, Panicker AJ, Dey AK, et al. (2024) Integrated approaches for immunotoxicity risk assessment: challenges and future directions. *Discover Toxicology* 1.
157. Yeku OO, Brentjens RJ (2016) Armored CAR T-cells: utilizing cytokines and pro-inflammatory ligands to enhance CAR T-cell anti-tumour efficacy. *Biochem Soc Trans* 44: 412-418.
158. Li X, Chen T, Li X, et al. (2024) Therapeutic targets of armored chimeric antigen receptor T cells navigating the tumor microenvironment. *Exp Hematol Oncol* 13: 96.
159. Han X, Wang Y, Wei J, et al. (2019) Multi-antigen-targeted chimeric antigen receptor T cells for cancer therapy. *J Hematol Oncol* 12: 128.
160. Singh AV, Bhardwaj P, Laux P, et al. (2024) AI and ML-based risk assessment of chemicals: predicting carcinogenic risk from chemical-induced genomic instability. *Front Toxicol* 6: 1461587.
161. Zitnik M, Li MM, Wells A, et al. (2024) Current and future directions in network biology. *Bioinform Adv* 4: vbae099.
162. Blay V, Pandiella A (2024) Strategies to boost antibody selectivity in oncology. *Trends Pharmacol Sci* 45: 1135-1149.
163. This S, Costantino S, Melichar HJ (2024) Machine learning predictions of T cell antigen specificity from intracellular calcium dynamics. *Sci Adv* 10: eadk2298.
164. Putignano G, Ruiperez-Campillo S, Yuan Z, et al. (2025) Mathematical models and computational approaches in CAR-T therapeutics. *Front Immunol* 16: 1581210.
165. McFaline-Figueroa JL, Srivatsan S, Hill AJ, et al. (2024) Multiplex single-cell chemical genomics reveals the kinase dependence of the response to targeted therapy. *Cell Genom* 4: 100487.
166. Biederstadt A, Manzar GS, Daher M (2022) Multiplexed engineering and precision gene editing in cellular immunotherapy. *Front Immunol* 13: 1063303.
167. Zambrano N, Froechlich G, Lazarevic D, et al. (2022) High-Throughput Monoclonal Antibody Discovery from Phage Libraries: Challenging the Current Preclinical Pipeline to Keep the Pace with the Increasing mAb Demand. *Cancers (Basel)* 14.
168. Lopez R, Regier J, Cole MB, et al. (2018) Deep generative modeling for single-cell transcriptomics. *Nat Methods* 15: 1053-1058.

169. Specht H, Emmott E, Petelski AA, et al. (2021) Single-cell proteomic and transcriptomic analysis of macrophage heterogeneity using SCoPE2. *Genome Biol* 22: 50.
170. Vijayan RSK, Kihlberg J, Cross JB, et al. (2022) Enhancing preclinical drug discovery with artificial intelligence. *Drug Discov Today* 27: 967-984.
171. Audagnotto M, Czechtizky W, De Maria L, et al. (2022) Machine learning/molecular dynamic protein structure prediction approach to investigate the protein conformational ensemble. *Sci Rep* 12: 10018.
172. Chen L, Balabanidou V, Remeta DP, et al. (2011) Structural instability tuning as a regulatory mechanism in protein-protein interactions. *Mol Cell* 44: 734-744.
173. Varadi M, Anyango S, Deshpande M, et al. (2022) AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic Acids Res* 50: D439-D444.
174. Maguire JB, Haddox HK, Strickland D, et al. (2021) Perturbing the energy landscape for improved packing during computational protein design. *Proteins* 89: 436-449.
175. Eastman P, Swails J, Chodera JD, et al. (2017) OpenMM 7: Rapid development of high performance algorithms for molecular dynamics. *PLoS Comput Biol* 13: e1005659.
176. Lyskov S, Gray JJ (2008) The RosettaDock server for local protein-protein docking. *Nucleic Acids Res* 36: W233-238.
177. Pratt OS, Elliott LG, Haon M, et al. (2025) AlphaFold 2, but not AlphaFold 3, predicts confident but unrealistic beta-solenoid structures for repeat proteins. *Comput Struct Biotechnol J* 27: 467-477.
178. Chang MP, Jin T, Gudinas AP, et al. (2025) Repetitive proteins that undergo large conformational changes evade structural prediction algorithms. *J Chem Phys* 163.
179. Dilchert J, Hofmann M, Unverdorben F, et al. (2022) Mammalian Display Platform for the Maturation of Bispecific TCR-Based Molecules. *Antibodies (Basel)* 11.
180. Smith SN, Harris DT, Kranz DM (2015) T Cell Receptor Engineering and Analysis Using the Yeast Display Platform. *Methods Mol Biol* 1319: 95-141.
181. Gai SA, Wittrup KD (2007) Yeast surface display for protein engineering and characterization. *Curr Opin Struct Biol* 17: 467-473.
182. Hartung T (2025) AI, agentic models and lab automation for scientific discovery - the beginning of scAInce. *Front Artif Intell* 8: 1649155.
183. Irvine DJ, Maus MV, Mooney DJ, et al. (2022) The future of engineered immune cell therapies. *Science* 378: 853-858.
184. Wang W, Hariharan M, Ding W, et al. (2025) Genetics and Environment Distinctively Shape the Human Immune Cell Epigenome. *bioRxiv*.
185. Kondilis-Mangum HD, Wade PA (2013) Epigenetics and the adaptive immune response. *Mol Aspects Med* 34: 813-825.
186. Sadria M, Layton A (2025) scVAEDer: integrating deep diffusion models and variational autoencoders for single-cell transcriptomics analysis. *Genome Biol* 26: 64.
187. Jung S (2025) Advances in modeling cellular state dynamics: integrating omics data and predictive techniques. *Anim Cells Syst (Seoul)* 29: 72-83.

188. Rodov A, Baniadam H, Zeiser R, et al. (2025) Towards the Next Generation of Data-Driven Therapeutics Using Spatially Resolved Single-Cell Technologies and Generative AI. *Eur J Immunol* 55: e202451234.
189. Ding J, Regev A (2021) Deep generative model embedding of single-cell RNA-Seq profiles on hyperspheres and hyperbolic spaces. *Nat Commun* 12: 2554.
190. Choi H, Kim H, Chung H, et al. (2025) Application of computational algorithms for single-cell RNA-seq and ATAC-seq in neurodegenerative diseases. *Brief Funct Genomics* 24.
191. Wang C, Liu ZP (2025) Diffusion-based generation of gene regulatory networks from scRNA-seq data with DigNet. *Genome Res* 35: 340-354.
192. Yeo GHT, Saksena SD, Gifford DK (2021) Generative modeling of single-cell time series with PRESCIENT enables prediction of cell trajectories with interventions. *Nat Commun* 12: 3222.
193. Weerakoon H, Mohamed A, Wong Y, et al. (2024) Integrative temporal multi-omics reveals uncoupling of transcriptome and proteome during human T cell activation. *NPJ Syst Biol Appl* 10: 21.
194. Wang X, Fan D, Yang Y, et al. (2023) Integrative multi-omics approaches to explore immune cell functions: Challenges and opportunities. *iScience* 26: 106359.
195. Singh AV, Bhardwaj P, Kishore V, et al. (2025) Multifunctional Nanoscale Pigments: Emerging Risks and Circular Strategies for a Sustainable Future. *Small Sci* 5: e202500240.
196. Lotfollahi M, Klimovskaia Susmelj A, De Donno C, et al. (2023) Predicting cellular responses to complex perturbations in high-throughput screens. *Mol Syst Biol* 19: e11517.
197. Lotfollahi M, Wolf FA, Theis FJ (2019) scGen predicts single-cell perturbation responses. *Nat Methods* 16: 715-721.
198. Metzner E, Southard KM, Norman TM (2025) Multiome Perturb-seq unlocks scalable discovery of integrated perturbation effects on the transcriptome and epigenome. *Cell Syst* 16: 101161.
199. Zhou P, Shi H, Huang H, et al. (2023) Single-cell CRISPR screens in vivo map T cell fate regulomes in cancer. *Nature* 624: 154-163.
200. Belk JA, Yao W, Ly N, et al. (2022) Genome-wide CRISPR screens of T cell exhaustion identify chromatin remodeling factors that limit T cell persistence. *Cancer Cell* 40: 768-786 e767.
201. Xia L, Komissarova A, Jacover A, et al. (2023) Systematic identification of gene combinations to target in innate immune cells to enhance T cell activation. *Nat Commun* 14: 6295.
202. Prochazka L, Michaels YS, Lau C, et al. (2022) Synthetic gene circuits for cell state detection and protein tuning in human pluripotent stem cells. *Mol Syst Biol* 18: e10886.
203. Sole R, Conde-Pueyo N, Pla-Mauri J, et al. (2024) Open problems in synthetic multicellularity. *NPJ Syst Biol Appl* 10: 151.

204. Nilsson A, Peters JM, Meimetis N, et al. (2022) Artificial neural networks enable genome-scale simulations of intracellular signaling. *Nat Commun* 13: 3069.
205. Racovita A, Jaramillo A (2020) Reinforcement learning in synthetic gene circuits. *Biochem Soc Trans* 48: 1637-1643.
206. Ding X, Zhang L, Fan M, et al. (2025) Network-based transfer of pan-cancer immunotherapy responses to guide breast cancer prognosis. *NPJ Syst Biol Appl* 11: 4.
207. DaSilva LF, Senan S, Patel ZM, et al. (2024) DNA-Diffusion: Leveraging Generative Models for Controlling Chromatin Accessibility and Gene Expression via Synthetic Regulatory Elements. *bioRxiv*.
208. Garcia BT, Westerfield L, Yelemali P, et al. (2025) Improving automated deep phenotyping through large language models using retrieval-augmented generation. *Genome Med* 17: 91.
209. Dong M, Wang L, Hu N, et al. (2025) Integration of multi-omics approaches in exploring intra-tumoral heterogeneity. *Cancer Cell Int* 25: 317.
210. van Lent P, Schmitz J, Abeel T (2023) Simulated Design-Build-Test-Learn Cycles for Consistent Comparison of Machine Learning Methods in Metabolic Engineering. *ACS Synth Biol* 12: 2588-2599.
211. Kitano S, Lin C, Foo JL, et al. (2023) Synthetic biology: Learning the way toward high-precision biological design. *PLoS Biol* 21: e3002116.
212. Daniszewski M, Crombie DE, Henderson R, et al. (2018) Automated Cell Culture Systems and Their Applications to Human Pluripotent Stem Cell Studies. *SLAS Technol* 23: 315-325.
213. Lei Y, Tsang JS (2025) Systems Human Immunology and AI: Immune Setpoint and Immune Health. *Annu Rev Immunol* 43: 693-722.
214. Gurdo N, Volke DC, McCloskey D, et al. (2023) Automating the design-build-test-learn cycle towards next-generation bacterial cell factories. *N Biotechnol* 74: 1-15.
215. Si Y, Zou J, Gao Y, et al. (2024) Foundation models in molecular biology. *Biophys Rep* 10: 135-151.
216. Moldwin A, Shehu A (2025) Foundation Models for AI-Enabled Biological Design. *arxiv*.
217. Wytock TP, Motter AE (2024) Cell reprogramming design by transfer learning of functional transcriptional networks. *Proc Natl Acad Sci U S A* 121: e2312942121.
218. Blaby IK, Cheng JF (2020) Building a custom high-throughput platform at the Joint Genome Institute for DNA construct design and assembly-present and future challenges. *Synth Biol (Oxf)* 5: ysaa023.
219. Ma Y, Zhang Z, Jia B, et al. (2024) Automated high-throughput DNA synthesis and assembly. *Heliyon* 10: e26967.
220. Rezalotfi A, Fritz L, Forster R, et al. (2022) Challenges of CRISPR-Based Gene Editing in Primary T Cells. *Int J Mol Sci* 23.
221. Kwak S, Lee H, Yu D, et al. (2025) Microfluidic Platforms for Ex Vivo and In Vivo Gene Therapy. *Biosensors (Basel)* 15.

222. Hunt AC, Rasor BJ, Seki K, et al. (2025) Cell-Free Gene Expression: Methods and Applications. *Chem Rev* 125: 91-149.
223. Lu Y (2024) The future of cell-free synthetic biology. *Biotechnol Notes* 5: A1-A3.
224. Craig T, Holland R, D'Amore R, et al. (2017) Leaf LIMS: A Flexible Laboratory Information Management System with a Synthetic Biology Focus. *ACS Synth Biol* 6: 2273-2280.
225. Zhou Y, Shao N, Bessa de Castro R, et al. (2020) Evaluation of Single-Cell Cytokine Secretion and Cell-Cell Interactions with a Hierarchical Loading Microwell Chip. *Cell Rep* 31: 107574.
226. Lim J, Park C, Kim M, et al. (2024) Advances in single-cell omics and multiomics for high-resolution molecular profiling. *Exp Mol Med* 56: 515-526.
227. Roth JP, Bajorath J (2024) Relationship between prediction accuracy and uncertainty in compound potency prediction using deep neural networks and control models. *Sci Rep* 14: 6536.
228. Hillson N, Caddick M, Cai Y, et al. (2019) Building a global alliance of biofoundries. *Nat Commun* 10: 2040.
229. Matzko R, Konur S (2024) Technologies for design-build-test-learn automation and computational modelling across the synthetic biology workflow: a review. *Network Modeling Analysis in Health Informatics and Bioinformatics* 13.
230. Goshisht MK (2024) Machine Learning and Deep Learning in Synthetic Biology: Key Architectures, Applications, and Challenges. *ACS Omega* 9: 9921-9945.
231. Holowko MB, Frow EK, Reid JC, et al. (2021) Building a biofoundry. *Synth Biol (Oxf)* 6: ysaa026.
232. Bultelle M, Casas A, Kitney R (2024) Engineering biology and automation-Replicability as a design principle. *Eng Biol* 8: 53-68.
233. Emmert-Streib F, Yli-Harja O (2022) What Is a Digital Twin? Experimental Design for a Data-Centric Machine Learning Perspective in Health. *Int J Mol Sci* 23.
234. McLaughlin JA, Beal J, Misirli G, et al. (2020) The Synthetic Biology Open Language (SBOL) Version 3: Simplified Data Exchange for Bioengineering. *Front Bioeng Biotechnol* 8: 1009.
235. Gierend K, Kruger F, Genehr S, et al. (2024) Provenance Information for Biomedical Data and Workflows: Scoping Review. *J Med Internet Res* 26: e51297.
236. Singh N, Lane S, Yu T, et al. (2025) A generalized platform for artificial intelligence-powered autonomous enzyme engineering. *Nat Commun* 16: 5648.
237. Rapp JT, Bremer BJ, Romero PA (2024) Self-driving laboratories to autonomously navigate the protein fitness landscape. *Nat Chem Eng* 1: 97-107.
238. Martin HG, Radivojevic T, Zucker J, et al. (2023) Perspectives for self-driving labs in synthetic biology. *Curr Opin Biotechnol* 79: 102881.
239. Melocchi A, Schmittlein B, Sadhu S, et al. (2025) Automated manufacturing of cell therapies. *J Control Release* 381: 113561.
240. Wang B, Chen RQ, Li J, et al. (2024) Interfacing data science with cell therapy manufacturing: where we are and where we need to be. *Cytotherapy* 26: 967-979.

241. Gemmati D, Longo G, Gallo I, et al. (2022) Host genetics impact on SARS-CoV-2 vaccine-induced immunoglobulin levels and dynamics: The role of TP53, ABO, APOE, ACE2, HLA-A, and CRP genes. *Front Genet* 13: 1028081.
242. Morgan RA, Chinnasamy N, Abate-Daga D, et al. (2013) Cancer regression and neurological toxicity following anti-MAGE-A3 TCR gene therapy. *J Immunother* 36: 133-151.
243. Birnbaum ME, Mendoza JL, Sethi DK, et al. (2014) Deconstructing the peptide-MHC specificity of T cell recognition. *Cell* 157: 1073-1087.
244. Cai Y, Chen R, Gao S, et al. (2022) Artificial intelligence applied in neoantigen identification facilitates personalized cancer immunotherapy. *Front Oncol* 12: 1054231.
245. Zou H, Liu W, Wang X, et al. (2024) Dynamic monitoring of circulating tumor DNA reveals outcomes and genomic alterations in patients with relapsed or refractory large B-cell lymphoma undergoing CAR T-cell therapy. *J Immunother Cancer* 12.
246. Bottini M, Ryu SJ, Terander AE, et al. (2025) The Ever-Evolving Regulatory Landscape Concerning Development and Clinical Application of Machine Intelligence: Practical Consequences for Spine Artificial Intelligence Research. *Neurospine* 22: 134-143.
247. Hersh W, Fultz Hollis K (2024) Results and implications for generative AI in a large introductory biomedical and health informatics course. *NPJ Digit Med* 7: 247.
248. Chao R, Mishra S, Si T, et al. (2017) Engineering biological systems using automated biofoundries. *Metab Eng* 42: 98-108.
249. Nettleton DF, Mari-Buye N, Marti-Soler H, et al. (2023) Smart Sensor Control and Monitoring of an Automated Cell Expansion Process. *Sensors (Basel)* 23.
250. Cheng F, Xie W, Zheng H (2024) Digital Twin Calibration for Biological System-of-Systems: Cell Culture Manufacturing Process. *arxiv*.
251. Hie BL, Shanker VR, Xu D, et al. (2024) Efficient evolution of human antibodies from general protein language models. *Nat Biotechnol* 42: 275-283.
252. Jonny, Sitepu EC, Nidom CA, et al. (2024) Ex Vivo-Generated Tolerogenic Dendritic Cells: Hope for a Definitive Therapy of Autoimmune Diseases. *Curr Issues Mol Biol* 46: 4035-4048.
253. Stucchi A, Maspes F, Montee-Rodrigues E, et al. (2024) Engineered Treg cells: The heir to the throne of immunotherapy. *J Autoimmun* 144: 102986.
254. Wang K, Song B, Zhu Y, et al. (2024) Peripheral nerve-derived CSF1 induces BMP2 expression in macrophages to promote nerve regeneration and wound healing. *NPJ Regen Med* 9: 35.
255. O'Brien H, Salm M, Morton LT, et al. (2024) A modular protein language modelling approach to immunogenicity prediction. *PLoS Comput Biol* 20: e1012511.
256. Alowidi N, Ali R, Sadaqah M, et al. (2024) Advancing Kidney Transplantation: A Machine Learning Approach to Enhance Donor-Recipient Matching. *Diagnostics (Basel)* 14.
257. Weimer ET, Newhall KA (2025) Machine learning enhanced immunologic risk assessments for solid organ transplantation. *Sci Rep* 15: 7943.

258. Alsaafeen BH, Ali BR, Elkord E (2025) Combinational therapeutic strategies to overcome resistance to immune checkpoint inhibitors. *Front Immunol* 16: 1546717.
259. Liu Z, Zhang J, Hong L, et al. (2025) Multiscale mathematical model-informed reinforcement learning optimizes combination treatment scheduling in glioblastoma evolution. *Sci Adv* 11: eadv3316.
260. Prelaj A, Miskovic V, Zanitti M, et al. (2024) Artificial intelligence for predictive biomarker discovery in immuno-oncology: a systematic review. *Ann Oncol* 35: 29-65.
261. Ahlquist KD, Sugden LA, Ramachandran S (2023) Enabling interpretable machine learning for biological data with reliability scores. *PLoS Comput Biol* 19: e1011175.
262. Shiferaw KB, Roloff M, Balaur I, et al. (2025) Guidelines and standard frameworks for artificial intelligence in medicine: a systematic review. *JAMIA Open* 8: ooae155.
263. Vithlani J, Hawksworth C, Elvidge J, et al. (2023) Economic evaluations of artificial intelligence-based healthcare interventions: a systematic literature review of best practices in their conduct and reporting. *Front Pharmacol* 14: 1220950.
264. Pannu J, Bloomfield D, MacKnight R, et al. (2025) Dual-use capabilities of concern of biological AI models. *PLoS Comput Biol* 21: e1012975.
265. Bloomfield D, Pannu J, Zhu AW, et al. (2024) AI and biosecurity: The need for governance. *Science* 385: 831-833.
266. Derraz B, Breda G, Kaempf C, et al. (2024) New regulatory thinking is needed for AI-based personalised drug and cell therapies in precision oncology. *NPJ Precis Oncol* 8: 23.
267. Groff-Vindman CS, Trump BD, Cummings CL, et al. (2025) The convergence of AI and synthetic biology: the looming deluge. *npj Biomedical Innovations* 2.
268. Undheim TA (2024) The whack-a-mole governance challenge for AI-enabled synthetic biology: literature review and emerging frameworks. *Front Bioeng Biotechnol* 12: 1359768.
269. Wu J, Koelzer VH (2024) Towards generative digital twins in biomedical research. *Comput Struct Biotechnol J* 23: 3481-3488.
270. Liu S, Guo LR (2024) Data Ownership in the AI-Powered Integrative Health Care Landscape. *JMIR Med Inform* 12: e57754.
271. Welzel C, Ostermann M, Smith HL, et al. (2025) Enabling secure and self determined health data sharing and consent management. *NPJ Digit Med* 8: 560.
272. Balsa-Canto E, Campo-Manzanares N, Moimenta AR, et al. (2025) Quantifying and managing uncertainty in systems biology: Mechanistic and data-driven models. *Current Opinion in Systems Biology* 42.
273. Zhou Z, Hu M, Salcedo M, et al. (2023) XAI meets Biology: A Comprehensive Review of Explainable AI in Bioinformatics Applications. *arxiv*.
274. Singh R, Paxton M, Auclair J (2025) Regulating the AI-enabled ecosystem for human therapeutics. *Commun Med (Lond)* 5: 181.
275. Niazi SK (2025) Regulatory Perspectives for AI/ML Implementation in Pharmaceutical GMP Environments. *Pharmaceuticals (Basel)* 18.
276. de Lima RC, Quaresma JAS (2025) Emerging technologies transforming the future of global biosecurity. *Front Digit Health* 7: 1622123.

277. Kumar D, Malin BA, Vishwanatha JK, et al. (2024) AI in Biomedicine-A Forward-Looking Perspective on Health Equity. *Int J Environ Res Public Health* 21.
278. Johnson KB, Horn IB, Horvitz E (2025) Pursuing Equity With Artificial Intelligence in Health Care. *JAMA Health Forum* 6: e245031.
279. Krenn M, Pollice R, Guo SY, et al. (2022) On scientific understanding with artificial intelligence. *Nat Rev Phys* 4: 761-769.
280. Alvarado R (2023) AI as an Epistemic Technology. *Sci Eng Ethics* 29: 32.