



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

Ola A Al-Ewaidat^{✉,1} 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

Article History

Received: November 06, 2025

Accepted: March 31, 2026

Published: May 08, 2026

Abstract

Recent advances in generative artificial intelligence have begun to reshape the field of cell and immune engineering. By learning the statistical and structural principles underlying biological systems, generative models can now design T-cell receptors, chimeric antigen receptors, and synthetic immune circuits that satisfy complex requirements for affinity, stability, and specificity. When integrated into automated design–build–test–learn pipelines, these models facilitate iterative cycles of hypothesis generation, experimental validation, and model refinement, thereby establishing a closed feedback loop between computational processes and biological systems. This review examines how AI-driven generative design is transforming immunoengineering across multiple scales, ranging from molecular recognition to cellular phenotype and clinical translation. It discusses the foundational architectures underlying generative modeling in biology, the emergence of adaptive biofoundries that connect digital design with manufacturing, and the translational pathways through which programmable immune cells may be advanced toward clinical application. The review also explores the ethical and regulatory dimensions of algorithmic biology, emphasizing the need for transparency, equitable access, and anticipatory governance. Collectively, these developments signal the emergence of a new paradigm—programmable immunity—in which biological design, therapeutic discovery, and ethical responsibility co-evolve within an integrated, 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 complex 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 constitutes a combinatorial experiment in molecular recognition, generated through stochastic recombination and refined by selective processes [1]. This distributed form of biological intelligence endows the immune system with exceptional specificity and plasticity; however, it also makes therapeutic design highly complex. Efforts to reprogram immune function, whether

through the engineering of monoclonal antibodies, the construction of chimeric antigen receptors (CARs), or the modulation of regulatory cell lineages, have historically relied on a combination of rational design and empirical approaches [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 cellular states is vast, multidimensional, and has been only sparsely sampled through experimental approaches [3,4].

* Corresponding Author:

Moawiah M. Naffaa, Independent Researcher, Mountain View, CA 94040, USA,
moawiahn@hotmail.com



© 2026 Copyright by the Authors.

Licensed as an open access article using a [CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/).

For most of the history of computational biology, artificial intelligence has been used primarily as an analytical and predictive tool. Discriminative and regression-based models have traditionally been trained to address questions such as “*What does this sequence do?*” or “*What structure will this sequence adopt?*” These models map existing biological sequences to inferred properties, thereby 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]. Within this framework, biology is treated primarily as an object to be explained rather than as a system to be actively designed or engineered.

Generative artificial intelligence introduces a fundamentally different epistemological framework. Rather than asking what a given biological sequence means, generative models pose the question of which sequences could exist to satisfy specified functional or structural objectives. This represents a shift from inference to invention, moving from the analysis of biological systems to the active design of novel biological forms. The central question therefore transitions from “What does this sequence do?” to “What sequence will perform this function?”. In this sense, generative AI transforms biology from a primarily descriptive science into a design-oriented discipline, in which learned statistical representations of evolution and biophysics are leveraged to create novel molecules, receptors, and cellular programs that do not exist in nature, as exemplified by generative protein language models and structure-generation systems [6,7].

In recent years, the emergence of generative artificial intelligence has introduced a qualitatively new paradigm in 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 revolutionized natural language and image generation—exemplified by models such as ESM-2, ProteinMPNN, and RFdiffusion—has been adapted to protein science, where it captures 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, these models give rise to what can be termed generative immunoengineering—the computational design of immune molecules and cellular programs guided by learned representations of underlying 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 antigenic contexts [12]. ProteinMPNN-based approaches have likewise 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 physically realistic conformational manifolds [6]. Beyond receptor design, multimodal generative frameworks are increasingly being developed to integrate transcriptomic, proteomic, and signaling data, enabling the prediction and, ultimately, the design of cellular phenotypes. These developments suggest that immune cells may soon be programmable entities, with their functional repertoires expanded through algorithmic design [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 historically been constrained by limited experimental throughput. It is now being replaced 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 increasingly functions as a feedback system, serving as an adaptive interface between computation and biology. This integration has already accelerated antibody discovery, improved TCR–peptide–HLA binding prediction, and enabled automated screening pipelines driven by active learning [16]. In the longer term, the coupling of *in silico* generation with *in vitro* verification may enable self-improving biomanufacturing ecosystems capable of producing tailored immune therapies with unprecedented speed and precision [17].

The opportunities presented by this paradigm are accompanied by significant challenges. The immune repertoire data used to train generative models remain incomplete and are often biased toward specific species, dis-

eases, and sequencing modalities, raising concerns regarding generalizability and the potential for off-target effects [18]. The interpretability of high-capacity models remains limited, and the regulatory frameworks governing AI-designed biologics are still in the early stages of development. 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 challenges will require the development of 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, with a particular focus on how generative algorithms are reshaping the design landscape of receptors, signaling modules, and cellular behaviors. The conceptual foundations are synthesized, recent technological advances are surveyed, and the translational and governance challenges accompany-

ing this accelerating field are outlined. The goal is to articulate a coherent framework for the 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 scope of therapeutic possibilities (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 that is 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 and empirical model grounded in target identification, scaffold selec-

tion, and iterative optimization. Breakthroughs such as monoclonal antibodies and first-generation CAR-T cells emerged from this approach [20]. However, dependence on sequential mutagenesis and experimental screening imposes significant constraints on both scale and efficiency. The theoretical diversity of immune receptor sequences

exceeds experimental capacity by many orders of magnitude [21], rendering comprehensive exploration of the antigen–receptor landscape unattainable and thereby creating a persistent bottleneck in immune engineering [22].

The introduction of generative artificial intelligence has begun to transform this landscape. Unlike discriminative algorithms, which classify existing data, generative models learn the statistical patterns underlying sequence–structure–function relationships and can therefore generate novel candidates that are consistent with these learned principles [23]. In this view, the immune system is no longer merely an object of analysis but a source of linguistic and structural priors from which algorithms infer the underlying 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 toward probabilistic generation [25].

Generative design is inherently shaped by the 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 generative capacity is constrained by the data on which they are trained, resulting in an inherent bias toward the exploitation of 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 overrepresent familiar motifs while underexploring rare yet potentially 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 may appear safe but ultimately constrain innovation [26,27].

Closed-loop design–build–test–learn frameworks partially mitigate this limitation by reintroducing experimental feedback, thereby expanding training distributions beyond historical datasets. However, even in adaptive pipelines, careful balance is required between exploiting known high-performing designs and exploring uncertain regions that may harbor breakthrough solutions. Recog-

nizing this exploration–exploitation trade-off is therefore essential for responsible generative immunoengineering, tempering optimism about algorithmic creativity with an awareness of its statistical and epistemic limitations.

Advances in protein foundation models provide the computational architecture for this transformation. Transformer-based language models such as ESM-2, ProtT5, and the 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, T-cell receptor (TCR), or antibody datasets, they learn both sequence syntax and the physicochemical principles governing 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 a study-specific application or fine-tuned adaptation of the open-source ProteinMPNN framework for 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 represents both a technological advancement and a 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 architectures, 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 is grounded in foundational advances in computational biology that have reshaped the representation of protein sequences, structures, and functions. 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 enabled transformer architectures, recurrent neural networks, and diffusion-based frameworks to learn 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 can capture latent biophysical properties, including secondary structure propensity, binding site probability, and thermostability [41, 42]. More recent architectures such as ESM-2 and the

MSA Transformer incorporate evolutionary information and multiple sequence alignments, achieving state-of-the-art performance in structure prediction and the inference of mutational effects [43–45]. 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 through discrete trial-and-error but by navigating a learned biological landscape [7,42].

Generative design occurs through steering within this latent space. Researchers do not directly manipulate sequences; instead, they guide 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 the latent space enables the generation of intermediate designs that smoothly trade-offs between properties such as affinity and stability, or activation and 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, as they are optimized within a potentially distorted representation of biological reality [48].

In generative immunoengineering, data quality dictates design outcomes. The clinical relevance, novelty, and safety of generated receptors or cellular states cannot surpass the informational limits of the datasets from which their latent representations 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-based frameworks such as ProteinMPNN, RFdiffusion, and Chroma model the conditional probability distribution of atomic arrangements conditioned on 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 generative models can propose receptor sequences that preserve structural fidelity while accommodating specific epitope geometries, a critical 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 sequences. 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 con-

textual constraints without requiring dense evolutionary alignments.

Physical constraints are incorporated through both architectural design and post-generation filtering. Structure-aware models such as ProteinMPNN and RFdiffusion represent proteins as graphs or geometric point clouds, where residues serve as 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 enables these models to internalize folding rules, hydrogen bonding patterns, and packing constraints implicitly through loss functions that penalize geometric deviations.

In immune applications, these physical constraints are further refined through 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 address 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 preserving the structural integrity of the conserved T cell receptor 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, thereby reducing the risk of generating biologically implausible de-

signs that are mathematically consistent but not physically viable.

To clarify how these architectures differ in data requirements, modeling strategies, 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 types of biological information on which they operate. Protein language models such as ESM-2 primarily learn from sequence statistics and are well suited for exploring diversity in T cell receptor 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 to ensure reproducibility and enable 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, several immune-specialized frameworks, such as PhysicoGPTCR and ProteinMPNN-TCR, represent research prototypes whose code availability varies across studies [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, such frameworks are explicitly designated as research prototypes rather than community tools. Generative CAR-design frameworks are fre-

quently developed within academic-industrial collaborations or commercial platforms and are often proprietary or study-specific [60,61]. In such cases, methodological details are provided in publications, whereas source code and trained models are not made publicly available. 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 cellular states, allowing prediction 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, in which molecular design and cellular 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, the models move beyond unconstrained generation toward constrained biological design aligned with therapeutic objectives.

The reliability of generative models critically depends on the quality, diversity, and annotation of the 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 enabling transparent benchmarking of model performance [72].

The theoretical strength of these models is complemented by their capacity for interpretability and embedding analysis, enabling mechanistic insight into learned representations. Techniques such as attention visualization, feature attribution, and latent-space interpolation have shown that protein language models implicitly capture biochemical hierarchies that resemble 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.

Together, these methodological pillars define 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 solely by human intuition 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 adaptive immune system's ability to distinguish self from non-self and to recognize an effectively unlimited

range of antigens arises from the extraordinary diversity of its receptor repertoire. 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 posed a challenge in 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 collectively contain millions of annotated T-cell and B-cell receptor sequences linked to antigenic or disease contexts [80–82]. 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 can 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 the understanding of immune recognition. High-resolution structural data from crystallography and cryo-electron microscopy, combined with molecular dynamics simulations, provide explicit insights into 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, enabling the 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 widely used strategy is transfer learning from large-scale protein corpora outside the immune system. 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 are then addressed through fine-tuning or conditioning on relatively small TCR–pMHC datasets, effectively leveraging general protein knowledge to mitigate the sparsity of immune-specific data.

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 structural information, they capture statistical regularities in 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 design–build–test–learn (DBTL) frameworks further help 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 enables models to progressively improve even when initial labeled datasets are limited.

Finally, weakly supervised and contrastive learning approaches enable 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 meth-

ods 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 enables models to generate receptors tailored to specific epitopes or immunological niches, rather than relying solely 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 with 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 [101–103]. 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 improves model fidelity and establishes an adaptive design framework that mirrors the iterative nature of immune evolution.

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 by developing standardized benchmarking platforms, curated cross-reactivity datasets, and transparent reporting frameworks for validation metrics that assess 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 underlying 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 integration of repertoire-scale learning with structural modeling thus constitutes the methodological core of generative immunoengineering, providing an analytical framework through which specificity can be understood, predicted, and engineered (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 boundaries 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 T-cell receptor (TCR) design, employ contextual embeddings that capture residue-level physicochemical properties while conditioning sequence generation on peptide–HLA complexes 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, including 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 solu-

tion. Reinforcement-learning based and active-learning frameworks address this challenge by optimizing multiple reward components simultaneously, identifying Pareto-optimal fronts that represent families of non-dominated receptor designs balancing competing objectives. Instead of reducing design to a single scalar score, these models preserve trade-offs among affinity, specificity, stability, and safety, enabling human designers to select candidates that best align with specific 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 [113–117].

5.1.1. Experimental Validation of AI-Designed TCRs

While many generative frameworks remain at the proof-of-concept stage, an increasing number of studies have demonstrated wet-lab validation of AI-designed or AI-optimized T-cell receptors (TCRs). These validations typically involve assessing surface expression using flow cytometry, antigen-specific binding assays, cytokine release measurements, and target-cell killing assays.

In one representative study, antigen-conditioned generative models were used to design novel T-cell receptor (TCR) sequences targeting defined peptide–HLA complexes. Selected candidates were expressed in primary T cells and demonstrated specific peptide–HLA binding, as confirmed by flow cytometry and multimer staining, indicating that the generated sequences formed functional surface-expressed 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 [120–122]. 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. Rather than optimizing solely for binding to a target peptide–HLA complex, these 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 and NetMHCpan-4.1 to estimate peptide presentation likelihood and potential immunogenicity, alongside deep-learning-based immunogenicity scoring frameworks such as DeepImmuno, which assess peptide–MHC immunogenic potential beyond binding affinity alone [126,127]. These predictors function as negative selection layers, ensuring that high-affinity designs are not advanced if they exhibit elevated risk of host immune responses. These layers collectively bias generation toward human-like sequences while removing candidates with elevated host-recognition risk.

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 using energy-based or steric-clash criteria, thereby reducing the risk of unintended tissue targeting.

Safety is further reinforced through physics-based and experimental filtering. Generated receptors are sub-

jected 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 T-cell receptor (TCR) engineering cannot be evaluated on 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, as well as root-mean-square deviation (RMSD) relative to known structures 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 studies indicate that AlphaFold-derived pipelines can produce near-native TCR–pMHC geometries in a subset of cases; however, performance remains variable and is often constrained by docking orientation errors and the intrinsic flexibility of CDR3 loops. Consequently, interface-focused metrics such as interface RMSD (iRMSD), DockQ, and contact recovery are frequently reported alongside global RMSD to provide a more accurate assessment of predictive quality [128,132]. For generative design, most studies do not yet report standardized head-to-head comparisons with homology

modeling or AlphaFold-Multimer on identical benchmark datasets. Instead, structure-aware generative pipelines typically use AlphaFold or AlphaFold-Multimer (or related predictors) as a screening and validation layer, reporting predicted interface RMSD (iRMSD), predicted aligned error (PAE)–based proxies, and docking scores, and subsequently prioritizing candidates for experimental binding and functional testing. The benchmarking practice can therefore be summarized as a two-stage evaluation: (i) computational assessment of interface geometry, including iRMSD, DockQ, and RMSD, along with confidence or PAE-based metrics; and (ii) experimental validation encompassing binding assays, functional activation, and target-cell killing assays [55].

5.1.3.1. Sequence- vs. Structure-Based Generation for TCR-Specific Design

Sequence-only language models (e.g., ESM-2–style approaches) are highly effective at learning repertoire priors and generating human-like, developable TCR variants that respect the statistical constraints of V/J gene 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- or interface-conditioned sequence design and diffusion-based geometric generation) are better suited for antigen-specific TCR engineering, as they can directly impose interface constraints such as contact geometry, shape complementarity, and steric feasibility, while localizing flexibility to CDR loops and 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 be-

tween 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 the selection of candidates for experimental validation. Together, these evaluation layers form a hierarchical metric system comprising sequence and structural 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 enhance 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 therefore propose residue substitutions that are likely to enhance affinity while avoiding steric clashes [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 enabling 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 generation 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 encoding 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 [153–155]. Such feedback loops have already produced CARs with modified hinge lengths and costimulatory 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 additional 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 transform CAR engineering from heuristic assembly into a rational optimization problem that can be addressed using machine learning.

5.3. Engineering Logic-Gated and Multiplexed Architectures

A further evolution of receptor design involves encoding 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 enhance selectivity and persistence while reducing 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 enable 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 relationships between module composition and functional synergy [165,166]. The ability to generate thousands of candidate architectures in silico and evaluate them using predictive scoring significantly accel-

erates 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 provides data for retraining generative models through active learning [168,169].

These advances have given rise to adaptive design–build–test–learn (DBTL) frameworks, in which computational generation and laboratory experimentation are linked through iterative feedback loops. 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 constitutes the operational core of generative immunoengineering.

Closed-loop pipelines are emerging in which model-generated receptor or construct libraries are automatically synthesized, expressed, and screened. The resulting activity and expression data are fed back 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 satisfy 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 design-build-test-learn cycle functions as the organizing principle linking algorithmic design to biological realization.

5.4.1. Computational Cost and Hardware Requirements

Generative immunoengineering pipelines exhibit substantial variation in computational cost, contingent upon the model class, scale, and stage of the workflow. Sequence-based protein language models, such as ESM-2 or immune-specific transformer architectures, can be fine-tuned and sampled on single high-memory GPUs or small multi-GPU servers. This computational profile makes them broadly accessible to many academic and translational research laboratories [7]. Inference for sequence generation typically requires timescales ranging from minutes to hours per batch, depending on model complexity, parameter size, and the selected sampling strategy.

Structure-aware and diffusion-based models, including ProteinMPNN and RFdiffusion, impose substantially higher computational demands [6,53]. Backbone-conditioned sequence design using ProteinMPNN is comparatively computationally efficient and can typically be executed on standard GPU workstations. In contrast, diffusion-based structure generation generally necessitates multi-GPU systems and longer execution times due to its iterative denoising process. Comprehensive three-dimensional generative workflows, when integrated with docking, energy minimization, and molecular dynamics refinement, typically require access to high-performance computing clusters or cloud-based GPU infrastructure.

However, these costs are typically concentrated in the early design phase. Once candidate receptors or constructs have been generated, downstream experimental screening becomes the dominant driver of both time and financial cost. 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].

Consequently, generative design is neither confined to supercomputing centers nor readily accessible to resource-constrained clinical laboratories. In practice, most translational pipelines are implemented through collaborations with academic computing centers, biofoundries, or cloud service providers, thereby enabling scalable computation without necessitating permanent in-house supercomputing infrastructure.

5.4.2. Hallucination Control and Physics-Based Filtering

Generative models may produce sequences or structures that are statistically plausible yet violate physical or biochemical constraints, a phenomenon commonly referred to as “hallucination.” Importantly, hallucination does not

refer solely to low-confidence structural predictions, such as those indicated by low pLDDT scores or high predicted aligned error. A more subtle and clinically relevant form involves high-confidence errors, wherein neural networks assign strong structural confidence to sequences that nevertheless fail to satisfy physical or thermodynamic constraints. These high-confidence hallucinations may satisfy learned statistical priors while still exhibiting physically implausible properties, such as exposed hydrophobic cores, a propensity for aggregation, destabilized secondary structure, or energetically unfavorable interfaces under physiological conditions.

At the structural level, generated sequences are typically evaluated using structure prediction or fold validation models such as AlphaFold or AlphaFold-Multimer to assess folding confidence and interface plausibility via metrics including 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 the risk of structural hallucination. Many workflows apply energy minimization and structural relaxation using force-field-based tools such as Rosetta, OpenMM, or GRO-MACS [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 while avoiding stable or promiscuous complexes with off-target peptides [176].

Molecular dynamics simulations provide an additional and particularly critical layer of validation for identifying 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 definitive filter for hallucinations. Generated receptors are subjected to early-stage expression screening, surface local-

ization assays, and binding validation experiments [179–181]. Candidates that fail to fold, traffic, or bind appropriately are eliminated before functional assays. This multi-layered strategy—comprising statistical generation, structural prediction, physics-based refinement, and experimental filtering—mitigates both low-confidence structural artifacts and high-confidence thermodynamic hallucinations. In doing so, it ensures that generative models do not operate as unconstrained proposal engines, but rather as integral components of a physically grounded and experimentally constrained 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 converge. This convergence marks the next stage of generative immunoengineering, in which molecular design and cellular behavior are jointly optimized within a unified, adaptive learning framework. 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 ability to engineer receptors and signaling modules has redefined the molecular architecture of immune cells; however, the next frontier of generative design extends beyond receptor composition toward 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 increasingly being adapted to capture these higher-order regulatory landscapes, offering a framework for proposing candidate strategies to influence cel-

lular 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 enable a 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 approximate a desired phenotype, such as enhanced persistence or reduced exhaustion. However, such phenotypic 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 persists between state interpolation and validated 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 occupy a high-dimensional state space defined by transcriptional, epigenetic, and metabolic variables that dynamically evolve 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, including 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 within, or sample from, this learned manifold to predict previously unseen or engineered phenotypes [186,187].

While these models are powerful in 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 intervention in causal regulatory networks, rather than 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 most are trained primarily on observational single-cell datasets rather than on 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 yield 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 [188–190]. These latent representations can be manipulated to simulate trajectories of transcriptional reprogramming. 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 analogue to Waddington’s epigenetic landscape that is 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 capture the coupled regulation of gene expression and metabolism in activated immune cells. Such models enable the generation of hypothetical phenotypes characterized by defined 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 transition of generative modeling from descriptive to prescriptive use involves linking 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 [198–200]. In macrophages and dendritic cells, similar models have been used to explore how combinations of signaling inputs reshape inflammatory versus tolerogenic polarization states. 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 cellular reprogramming as an optimization problem, in which artificial intelligence can assist in identifying promising sequences of interven-

tions, enabling dynamic control of gene regulatory networks rather than static modification of individual targets.

6.3. Linking Generative Models to Synthetic Circuits

Generative modeling also provides a computational substrate for the design of synthetic gene circuits capable of driving desired cell state transitions [202,203]. Once a target phenotype, such as resistance to exhaustion, enhanced persistence, or altered cytokine balance, is defined in latent space, AI models can identify candidate regulatory motifs or signaling pathways that may 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 the mapping between sequence composition and expression amplitude or inducibility, generative models can design synthetic regulatory elements that enable precise transcriptional tuning in 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 enables real-time assessment of how engi-

neered interventions reshape cellular states [209]. Data from each iteration are used to update generative priors, improving accuracy and adaptability.

This closed-loop paradigm parallels the design–build–test–learn cycles established in molecular engineering, but operates at the systems level of cellular behavior [210,211]. Extending the molecular DBTL (design–build–test–learn) framework to the cellular systems level enables iterative refinement of both molecular components and emergent phenotypes within a unified feedback architecture. In these frameworks, the model acts as a control algorithm that continuously adjusts interventions to maintain desired phenotypic states, requiring standardized data formats and interoperable protocols to ensure reproducibility and safe automation. These systems could ultimately form the basis of autonomous, adaptive immunoengineering platforms in which AI proposes genetic or pharmacological modifications, laboratory systems execute them through automated microfluidic experimentation, and the resulting data are used to continuously retrain and refine 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 increasingly 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 also 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, the DBTL (design–build–test–learn) framework transforms immunoengineering from a sequential workflow into a continuously adaptive system that integrates discovery, validation, and optimization into a unified process.

7.1. The Design Phase: Generative Hypothesis Formation

In the generative paradigm, design constitutes a computational experiment in which the model explores the proba-

bility 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, including affinity, stability, specificity, metabolic resilience, and manufacturability [215,216].

Multi-objective reinforcement learning (MORL) and Pareto-front optimization are increasingly employed to balance these criteria, ensuring that improvements in one property do not come at the expense of others. 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 and guide exploration toward regions of the design space where model 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 burden while improving 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 systems. 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 synthetic constructs into cellular systems. CRISPR/Cas systems, base editing, and transposon-mediated delivery methods enable targeted insertion of designed sequences into immune cell genomes, often at safe harbor loci that support stable and 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, enabling a traceable linkage between computational proposals and their 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 layer of validation between computational design and cellular implementation. 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 associated with testing AI-generated sequences. Integration with laboratory information management systems ensures that metadata, sequence provenance, and performance metrics are seamlessly fed back into the digital design environment [224].

7.3. The Test Phase: High-Dimensional and Multiscale Evaluation

Testing constitutes the sensory layer of the design–build–test–learn system, translating experimental outcomes into quantitative metrics that inform subsequent 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 advances in multi-omic screening have enhanced the resolution and depth of analytical 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 using unsupervised embedding techniques and graph-based clustering methods to extract latent features that represent functional archetypes. Statistical coupling between design parameters and phenotypic readouts allows causal inference about which molecular features drive performance. Importantly, uncertainty quantification metrics inform the selection of candidates for in-depth mechanistic analysis, ensuring that the testing phase not only validates designs but also enhances 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 constrain the learning rate of the overall generative system. As a consequence, advances in generative immunoengineering are intrinsically coupled with 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 learning phase closes the feedback loop by translating 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 identify experiments that 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, whereby successful experimental outcomes generate reward signals that bias subsequent generative sampling toward more productive directions.

In practice, these adaptive feedback mechanisms give rise to a cyber-physical learning system—an integrated framework 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 a substantial 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, comprising variation, selection, and retention implemented within an AI-governed experimental ecosystem.

7.5. Automation, Data Infrastructure, and Self-Optimizing Biofoundries

At scale, DBTL becomes inseparable from automation. Modern biofoundries integrate 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, including temperature, reagent usage, cell viability, and expression levels, are streamed to cloud infrastructures, where they 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 instantane-

ous 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 as 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 dataset functions as an institutional memory, enabling transfer learning across projects and facilitating 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 biological knowledge production. 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 transforms immunoengineering into a continuously learning infrastructure. Each iteration expands the collective intelligence encoded in both digital models and biological systems, gradually approaching a regime in which the boundaries between

designing, testing, and learning from immunity become increasingly 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 epistemological paradigm for biological design, in which 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, evaluated, 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 the approval of individual entities toward the 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 document design lineage, training data provenance, conditioning variables, and objective functions across iter-

ative development cycles. 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 the mechanism of action, state-space occupancy, or risk profile would necessitate additional preclinical or clinical evaluation.

This challenge is further amplified by the biological reality that immune cells occupy a high-dimensional state space defined by transcriptional, epigenetic, and metabolic variables, all of which 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 the absence of hazardous 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 prob-

abilistic 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 weeks, although this 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 enables *adaptive therapy*, in which continuous molecular monitoring through circulating tumor DNA, antigen-escape profiling, or cytokine dynamics is fed back into retraining models to update therapeutic 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 dynamic 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, enabling 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 dis-

cipline, 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 platforms directly to robotic assembly systems, viral vector packaging workflows, and closed-system cell expansion, creating an unbroken digital thread from algorithmic design to physical manufacture. Each construct carries a persistent identifier that links its computational origin to its production batch, ensuring traceability throughout 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 provide continuous data streams to these control layers, enabling 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 reconceptualize Good Manufacturing Practice (GMP) as a dynamic rather than static framework, in which quality is ensured through continuous sensing, prediction, and correction, rather than reliance on 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 enhance selectivity and persistence while mitigating off-target 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 inducing systemic immunosuppression. Regenerative medicine may benefit from engineered macrophage or tolerogenic antigen-presenting cell designs optimized for cytokine balance and metabolic resilience, thereby 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, ensuring that overreliance on automated predictions is avoided 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, thereby reducing logistical complexity and dependence on cold-chain systems. 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 preventing 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 requires explicit consent frameworks that define the use of secondary and longitudinal data. Questions of intellectual property, including 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. Safeguard-

ing measures analogous to existing biosecurity treaties will be essential, along with equitable access to computational infrastructure to prevent the concentration of capability within a small number of 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 this horizon lies the prospect of a continuously learning therapeutic ecosystem in which each patient outcome refines the generative models that inform subsequent design. In this envisioned future, therapeutic innovation and biological understanding become increasingly inseparable, and the immune system is reconceptualized as a programmable interface between computation and living systems. This is 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 are 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 governance and oversight frameworks required to ensure their safe and equitable application [267,268]. This dual responsibility to engineer biology and to define the ethics that govern it characterizes an emerging field of genera-

tive bioethics that is 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 post hoc to 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 reinforcement learning reward signals, ensuring that generated receptors or cell programs do not systematically favor 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 modeling 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 are not achieved 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, enabling models to adaptively constrain exploration of designs associated with elevated immunogenicity, off-target activity, or misuse potential. In this framework, ethics is not treated as a static compliance checkpoint but as 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 gen-

erative 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 regarding 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 [269–271]. In clinical settings, such adaptive consent mechanisms must be coupled with continuous feedback loops between patients and therapeutic models, ensuring that participants retain agency as their biological and computational 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. Incorporating underrepresented populations into training datasets is also essential to mitigate demographic bias and to ensure the 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. Recent advances in attention mapping, saliency analysis, and uncertainty quantification have enhanced interpretability of how models generate biological designs; however, a persistent epistemic gap remains between statistical pattern recognition and causal biological reasoning. 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 regarded as an integral component of 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. In the absence of such transparency, the reproducibility crisis observed in other disciplines may extend to synthetic immunology, thereby undermining confidence in AI-mediated design.

Regulatory institutions now face the challenge of governing entities that are not static products, but dynamic, continuously 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 Good Manufacturing Practice (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 proliferate across open scientific ecosystems, the distinction between beneficial and potentially hazardous applications becomes increasingly permeable. Consequently, international biosecurity regimes—traditionally centered on material agents and physical laboratory environments—must expand to incorporate informational biosecurity, encompassing the governance of digital models, codebases, and data pipelines capable of generating biologically relevant 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 integrate expertise in AI safety and cybersecurity alongside traditional biosafety expertise.

The societal implications of generative immunoengineering extend beyond bioethics into the political economy of biomedical innovation. Given that generative design relies on substantial computational infrastructure and frequently proprietary datasets, there is a concomitant risk that both technological capability and, by extension, access to therapeutic innovation may become increasingly concentrated within a small number of insti-

tutions 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 only a matter of distributive justice but also of scientific robustness. Diversity in training data improves model generalizability and reduces bias, thereby enhancing the safety and effectiveness of therapies across global populations. 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 profound. Generative immunoengineering reconfigures the relationship between hypothesis and experimentation, transforming scientific 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 understanding the underlying causal grammar, the locus of scientific agency shifts from the individual human investigator to a hybrid human-machine system. In this configuration, agency becomes distributed as 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, integrating generative immunoengineering into clinical and societal systems will require new governance frameworks that are anticipatory rather than reactive in nature. 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 regarding 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 hu-

man responsibility for defining its boundaries and purposes becomes more, not less, essential. The future of programmable immunity will therefore depend not only on the sophistication of its algorithms, but also on the moral architecture of the institutions that steward them [268]. Only through ethically adaptive governance can programmable immunity mature into a discipline that preserves both biological integrity and public trust in the era of generative medicine.

10. Conclusions

The convergence of generative artificial intelligence, cellular engineering, and immunology marks a transformative moment in the life sciences. What began as an effort to improve receptor design has evolved into a broader framework for understanding and engineering biological systems themselves. 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 reframes biology as an editable and self-informing system, in which learning becomes a property of both the organism and the methodologies used to investigate 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 enhances the system’s intelligence by converting empirical outcomes into computational insight, thereby enabling the co-evolution of therapeutic design and biological understanding. In this model, medicine becomes a learning enterprise, guided by algorithms that refine their predictions through every patient and experiment. Such integration signals the emergence of a “living laboratory” paradigm, in which computation, experimentation, and clinical feedback function as an integrated, self-optimizing network.

From a translational perspective, this paradigm reshapes both the structure and tempo of biomedical innovation. The capacity to generate, validate, and deploy patient-specific immune receptors within compressed timeframes enables precision immunotherapy to respond dynamically to both individual- and population-level variation. 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 term, this architectural paradigm may extend beyond immunology to integrate generative genomics, regenerative medicine, and neural interface design into a unified framework for 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 sustaining a balance between innovation and accountability, openness and security, and 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 calls for a reconceptualization of how biological intervention is understood and approached. 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, while remaining grounded in 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 more effective immune cells, but also to develop more responsible scientific, ethical, and societal systems through which such capabilities can be guided toward the enduring goal of human well-being.

List of Abbreviations

AI	Artificial Intelligence
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
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.A., M.M.N.; Methodology: M.M.N.; Formal analysis: O.A.A., M.M.N.; Investigation: M.M.N.; Resources: M.M.N.; Writing—original draft preparation: O.A.A., M.M.N.; Writing—review and editing: O.A.A., M.M.N.; Visualization: M.M.N.; Supervision: M.M.N.; Project administration: O.A.A., M.M.N. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest

The authors declare no conflicts of interest.

Funding

The study did not receive any external funding and was conducted using only institutional resources.

Acknowledgments

Declared 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.M.N. for this manuscript, with the assistance of AI-based visualization tools.

Supplementary Materials

Supplementary material associated with this article has been published online and is available at: <https://doi.org/10.69709/CTEC.2026.126651>.

References

- [1] Quiroz, R.N.; Camacho, J.V.; Peña, E.Z.; Lemus, Y.B.; López-Fernández, C.; Escorcia, L.G.; Fernández-Ponce, C.; Cobos, M.R.; Moreno, J.F.; Fiorillo-Moreno, O.; et al. Multiscale Information Processing in the Immune System. *Front. Immunol.* **2025**, *16*, 1563992. [CrossRef]
- [2] Dewaker, V.; Morya, V.K.; Kim, Y.H.; Park, S.T.; Kim, H.S.; Koh, Y.H. Revolutionizing Oncology: The Role of Artificial Intelligence (AI) as an Antibody Design, and Optimization Tools. *Biomark. Res.* **2025**, *13*, 52. [CrossRef]
- [3] Rafelski, S.M.; Theriot, J.A. Establishing a Conceptual Framework for Holistic Cell States and State Transitions. *Cell* **2024**, *187*, 2633–2651. [CrossRef] [PubMed]
- [4] Leary, A.Y.; Scott, D.; Gupta, N.T.; Waite, J.C.; Skokos, D.; Atwal, G.S.; Hawkins, P.G. Designing Meaningful Continuous Representations of T Cell Receptor Sequences with Deep Generative Models. *Nat. Commun.* **2024**, *15*, 4271. [CrossRef]
- [5] Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Židek, A.; Potapenko, A.; et al. Highly Accurate Protein Structure Prediction with AlphaFold. *Nature* **2021**, *596*, 583–589. [CrossRef] [PubMed]
- [6] Dauparas, J.; Anishchenko, I.; Bennett, N.; Bai, H.; Ragotte, R.J.; Milles, L.F.; Wicky, B.I.M.; Courbet, A.; de Haas, R.J.; Bethel, N.; et al. Robust Deep Learning–Based Protein Sequence Design Using ProteinMPNN. *Science* **2022**, *378*, 49–56. [CrossRef]
- [7] Rives, A.; Meier, J.; Sercu, T.; Goyal, S.; Lin, Z.; Liu, J.; Guo, D.; Ott, M.; Zitnick, C.L.; Ma, J.; et al. Biological Structure and Function Emerge from Scaling Unsupervised Learning to 250 Million Protein Sequences. *Proc. Natl. Acad. Sci. USA* **2021**, *118*, e2016239118. [CrossRef]
- [8] Ibrahim, M.; Al Khalil, Y.; Amirrajab, S.; Sun, C.; Breeuwer, M.; Pluim, J.; Elen, B.; Ertaylan, G.; Dumontier, M. Generative AI for Synthetic Data across Multiple Medical Modalities: A Systematic Review of Recent Developments and Challenges. *Comput. Biol. Med.* **2025**, *189*, 109834. [CrossRef]
- [9] Zhang, P.; Wei, L.; Li, J.; Wang, X. Artificial Intelligence-Guided Strategies for Next-Generation Biological Sequence Design. *Natl. Sci. Rev.* **2024**, *11*, nwae343. [CrossRef]
- [10] Ertel, M.; Moretti, R.; Meiler, J.; Schoeder, C.T. Self-Supervised Machine Learning Methods for Protein Design Improve Sampling but Not the Identification of High-Fitness Variants. *Sci. Adv.* **2025**, *11*, eadr7338. [CrossRef] [PubMed]
- [11] Krapp, L.F.; Meireles, F.A.; Abriata, L.A.; Devilard, J.; Vacle, S.; Marcaida, M.J.; Peraro, M.D. Context-Aware Geometric Deep Learning for Protein Sequence Design. *Nat. Commun.* **2024**, *15*, 6273. [CrossRef]
- [12] Ma, J.; Li, H.; Hu, Y.F.; Huang, J.D. Physicochemically Informed Dual-Conditioned Generative Model of T-Cell Receptor Variable Regions for Cellular Therapy. *arXiv* **2025**, arXiv:2510.05747. [CrossRef]
- [13] Chouleur, T.; Etchegaray, C.; Villain, L.; Lesur, A.; Ferté, T.; Rossi, M.; Andrique, L.; Simoncini, C.; Giacobbi, A.; Gambaretti, M.; et al. 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* **2025**, *157*, 573–587. [CrossRef]
- [14] Berson, E.; Chung, P.; Espinosa, C.; Montine, T.J.; Aghaeepour, N. Unlocking Human Immune System Complexity through AI. *Nat. Methods* **2024**, *21*, 1400–1402. [CrossRef]
- [15] Gao, S.; Fang, A.; Huang, Y.; Giunchiglia, V.; Noori, A.; Schwarz, J.R.; Ektefaie, Y.; Kondic, J.; Zitnik, M. Empowering Biomedical Discovery with AI Agents. *Cell* **2024**, *187*, 6125–6151. [CrossRef]
- [16] O'Donnell, T.J.; Kanduri, C.; Isacchini, G.; Limenitakis, J.P.; Brachman, R.A.; Alvarez, R.A.; Haff, I.H.; Sandve, G.K.; Greiff, V. Reading the Repertoire: Progress in Adaptive Immune Receptor Analysis Using Machine Learning. *Cell Syst.* **2024**, *15*, 1168–1189. [CrossRef] [PubMed]
- [17] Liu, Y.; Zhang, L.; Jiang, Z.; Tian, X.; Li, P.; Wu, P.; Du, W.; Yuan, B.; Xie, C.; Bu, G.; et al. Applications of Artificial Intelligence in Biotech Drug Discovery and Product Development. *Medcomm* **2025**, *6*, e70317. [CrossRef]
- [18] Katoh, H.; Komura, D.; Furuya, G.; Ishikawa, S. Immune Repertoire Profiling for Disease Pathobiology. *Pathol. Int.* **2023**, *73*, 1–11. [CrossRef]

- [19] Ou, Y.; Guo, S. Safety Risks and Ethical Governance of Biomedical Applications of Synthetic Biology. *Front. Bioeng. Biotechnol.* **2023**, *11*, 1292029. [CrossRef] [PubMed]
- [20] Kong, Y.; Li, J.; Zhao, X.; Wu, Y.; Chen, L. CAR-T Cell Therapy: Developments, Challenges and Expanded Applications from Cancer to Autoimmunity. *Front. Immunol.* **2024**, *15*, 1519671. [CrossRef]
- [21] de Greef, P.C.; Oakes, T.; Gerritsen, B.; Ismail, M.; Heather, J.M.; Hermesen, R.; Chain, B.; de Boer, R.J. The Naive T-Cell Receptor Repertoire Has an Extremely Broad Distribution of Clone Sizes. *eLife* **2020**, *9*, e49900. [CrossRef] [PubMed]
- [22] Tshouridis, O.; Xu, M.; Song, F.; Savoldo, B.; Dotti, G. The Landscape of CAR-Engineered Innate Immune Cells for Cancer Immunotherapy. *Nat. Cancer* **2025**, *6*, 1145–1156. [CrossRef]
- [23] Gangwal, A.; Ansari, A.; Ahmad, I.; Azad, A.K.; Kumarasamy, V.; Subramanian, V.; Wong, L.S. Generative Artificial Intelligence in Drug Discovery: Basic Framework, Recent Advances, Challenges, and Opportunities. *Front. Pharmacol.* **2024**, *15*, 1331062. [CrossRef]
- [24] Johnson, J.A.; Bergman, D.R.; Rocha, H.L.; Zhou, D.L.; Cramer, E.; Mclean, I.C.; Dance, Y.W.; Booth, M.; Nicholas, Z.; Lopez-Vidal, T.; et al. Human Interpretable Grammar Encodes Multicellular Systems Biology Models to Democratize Virtual Cell Laboratories. *Cell* **2025**, *188*, 4711–4733. [CrossRef]
- [25] Weber, C.R.; Rubio, T.; Wang, L.; Zhang, W.; Robert, P.A.; Akbar, R.; Snapkov, I.; Wu, J.; Kuijjer, M.L.; Tarazona, S.; et al. Reference-Based Comparison of Adaptive Immune Receptor Repertoires. *Cell Rep. Methods* **2022**, *2*, 100269. [CrossRef]
- [26] Brown, N.; Fiscato, M.; Segler, M.H.; Vaucher, A.C. GuacaMol: Benchmarking Models for De Novo Molecular Design. *J. Chem. Inf. Model.* **2019**, *59*, 1096–1108. [CrossRef] [PubMed]
- [27] Sanchez-Lengeling, B.; Aspuru-Guzik, A. Inverse Molecular Design Using Machine Learning: Generative Models for Matter Engineering. *Science* **2018**, *361*, 360–365. [CrossRef] [PubMed]
- [28] Bjerregaard, A.; Groth, P.M.; Hauberg, S.; Krogh, A.; Boomsma, W. Foundation Models of Protein Sequences: A Brief Overview. *Curr. Opin. Struct. Biol.* **2025**, *91*, 103004. [CrossRef]
- [29] Fox, D.R.; Taveneau, C.; Clement, J.; Grinter, R.; Knott, G.J. Code to Complex: AI-Driven De Novo Binder Design. *Structure* **2025**, *33*, 1631–1642. [CrossRef]
- [30] Wang, M.; Patsenker, J.; Li, H.; Kluger, Y.; Kleinstein, S.H. Supervised Fine-Tuning of Pre-Trained Antibody Language Models Improves Antigen Specificity Prediction. *PLoS Comput. Biol.* **2025**, *21*, e1012153. [CrossRef]
- [31] Chen, L.T.; Quinn, Z.; Dumas, M.; Peng, C.; Hong, L.; Lopez-Gonzalez, M.; Mestre, A.; Watson, R.; Vincoff, S.; Zhao, L.; et al. Target Sequence-Conditioned Design of Peptide Binders Using Masked Language Modeling. *Nat. Biotechnol.* **2025**. [CrossRef]
- [32] Gasser, H.-C.; Rajan, A.; Alfaro, J.A. A Novel Decoding Strategy for ProteinMPNN to Design with Less Visibility to Cytotoxic T-Lymphocytes. *Comput. Struct. Biotechnol. J.* **2025**, *27*, 3693–3703. [CrossRef] [PubMed]
- [33] Mei, A.; Letscher, K.P.; Reddy, S. 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.* **2024**, *90*, 103223. [CrossRef]
- [34] Zhang, J.; Che, Y.; Liu, R.; Wang, Z.; Liu, W. Deep Learning-Driven Multi-Omics Analysis: Enhancing Cancer Diagnostics and Therapeutics. *Brief. Bioinform.* **2025**, *26*, bbaf440. [CrossRef]
- [35] Sandberg, T.E.; Salazar, M.J.; Weng, L.L.; Palsson, B.O.; Feist, A.M. The Emergence of Adaptive Laboratory Evolution as an Efficient Tool for Biological Discovery and Industrial Biotechnology. *Metab. Eng.* **2019**, *56*, 1–16. [CrossRef]
- [36] Zhu, L.; Mou, W.; Hong, C.; Yang, T.; Lai, Y.; Qi, C.; Lin, A.; Zhang, J.; Luo, P. The Evaluation of Generative AI Should Include Repetition to Assess Stability. *JMIR Mhealth Uhealth* **2024**, *12*, e57978. [CrossRef]
- [37] Mirakhori, F.; Niazi, S.K. Harnessing the AI/ML in Drug and Biological Products Discovery and Development: The Regulatory Perspective. *Pharmaceuticals* **2025**, *18*, 47. [CrossRef]
- [38] Teng, F.; Cui, T.; Zhou, L.; Gao, Q.; Zhou, Q.; Li, W. Programmable Synthetic Receptors: The Next-Generation of Cell and Gene Therapies. *Signal Transduct. Target. Ther.* **2024**, *9*, 7. [CrossRef]
- [39] Weissenow, K.; Rost, B. Are Protein Language Models the New Universal Key? *Curr. Opin. Struct. Biol.* **2025**, *91*, 102997. [CrossRef] [PubMed]
- [40] Chen, Y.; Wang, Z.; Zeng, X.; Li, Y.; Li, P.; Ye, X.; Sakurai, T. Molecular Language Models: RNNs or Transformer? *Brief. Funct. Genom.* **2023**, *22*, 392–400. [CrossRef]
- [41] Brandes, N.; Ofer, D.; Peleg, Y.; Rappoport, N.; Linial, M. ProteinBERT: A Universal Deep-Learning Model of Protein Sequence and Function. *Bioinformatics* **2022**, *38*, 2102–2110. [CrossRef] [PubMed]
- [42] Alley, E.C.; Khimulya, G.; Biswas, S.; AlQuraishi, M.; Church, G.M. Unified Rational Protein Engineering with Sequence-Based Deep Representation Learning. *Nat. Methods* **2019**, *16*, 1315–1322. [CrossRef]
- [43] Lupo, U.; Sgarbossa, D.; Bitbol, A.-F. Protein Language Models Trained on Multiple Sequence Alignments Learn Phylogenetic Relationships. *Nat. Commun.* **2022**, *13*, 6298. [CrossRef]
- [44] Lin, Z.; Akin, H.; Rao, R.; Hie, B.; Zhu, Z.; Lu, W.; Smetanin, N.; Verkuil, R.; Kabeli, O.; Shmueli, Y.; et al. Evolutionary-Scale Prediction of Atomic-Level Protein Structure with a Language Model. *Science* **2023**, *379*, 1123–1130. [CrossRef]

- [45] Chen, Y.; Xu, Y.; Liu, D.; Xing, Y.; Gong, H. An End-to-End Framework for the Prediction of Protein Structure and Fitness from Single Sequence. *Nat. Commun.* **2024**, *15*, 7400. [\[CrossRef\]](#)
- [46] Wang, M.; Patsenker, J.; Li, H.; Kluger, Y.; Kleinstein, S.H. Language Model-Based B Cell Receptor Sequence Embeddings Can Effectively Encode Receptor Specificity. *Nucleic Acids Res.* **2024**, *52*, 548–557. [\[CrossRef\]](#) [\[PubMed\]](#)
- [47] Ornes, S. Researchers Turn to Deep Learning to Decode Protein Structures. *Proc. Natl. Acad. Sci. USA* **2022**, *119*, e2202107119. [\[CrossRef\]](#) [\[PubMed\]](#)
- [48] Heil, B.J.; Hoffman, M.M.; Markowitz, F.; Lee, S.-I.; Greene, C.S.; Hicks, S.C. Reproducibility Standards for Machine Learning in the Life Sciences. *Nat. Methods* **2021**, *18*, 1132–1135. [\[CrossRef\]](#) [\[PubMed\]](#)
- [49] Justyna, M.; Zirbel, C.; Antczak, M.; Szachniuk, M. Graph Neural Network and Diffusion Model for Modeling RNA Interatomic Interactions. *Bioinformatics* **2025**, *41*, btaf515. [\[CrossRef\]](#)
- [50] Soleymani, F.; Paquet, E.; Viktor, H.L.; Michalowski, W. Structure-Based Protein and Small Molecule Generation Using EGNN and Diffusion Models: A Comprehensive Review. *Comput. Struct. Biotechnol. J.* **2024**, *23*, 2779–2797. [\[CrossRef\]](#)
- [51] Baek, M.; DiMaio, F.; Anishchenko, I.; Dauparas, J.; Ovchinnikov, S.; Lee, G.R.; Wang, J.; Cong, Q.; Kinch, L.N.; Schaeffer, R.D.; et al. Accurate Prediction of Protein Structures and Interactions Using a Three-Track Neural Network. *Science* **2021**, *373*, 871–876. [\[CrossRef\]](#)
- [52] Wu, K.E.; Yang, K.K.; Berg, R.v.D.; Alamdari, S.; Zou, J.Y.; Lu, A.X.; Amini, A.P. Protein Structure Generation via Folding Diffusion. *Nat. Commun.* **2024**, *15*, 1059. [\[CrossRef\]](#)
- [53] Watson, J.L.; Juergens, D.; Bennett, N.R.; Trippe, B.L.; Yim, J.; Eisenach, H.E.; Ahern, W.; Borst, A.J.; Ragotte, R.J.; Milles, L.F.; et al. De Novo Design of Protein Structure and Function with RFdiffusion. *Nature* **2023**, *620*, 1089–1100. [\[CrossRef\]](#)
- [54] Fernández-Quintero, M.L.; Pomarici, N.D.; Loeffler, J.R.; Seidler, C.A.; Liedl, K.R. T-Cell Receptor CDR3 Loop Conformations in Solution Shift the Relative V α -V β Domain Distributions. *Front. Immunol.* **2020**, *11*, 1440. [\[CrossRef\]](#)
- [55] Quast, N.P.; Abanades, B.; Guloglu, B.; Karuppiiah, V.; Harper, S.; Raybould, M.I.J.; Deane, C.M. T-Cell Receptor Structures and Predictive Models Reveal Comparable Alpha and Beta Chain Structural Diversity Despite Differing Genetic Complexity. *Commun. Biol.* **2025**, *8*, 362. [\[CrossRef\]](#)
- [56] Bennett, N.R.; Coventry, B.; Goresnik, I.; Huang, B.; Allen, A.; Vafeados, D.; Peng, Y.P.; Dauparas, J.; Baek, M.; Stewart, L.; et al. Improving De Novo Protein Binder Design with Deep Learning. *Nat. Commun.* **2023**, *14*, 2625. [\[CrossRef\]](#)
- [57] Cao, L.; Coventry, B.; Goresnik, I.; Huang, B.; Sheffler, W.; Park, J.S.; Jude, K.M.; Marković, I.; Kadam, R.U.; Verschuere, K.H.G.; et al. Design of Protein-Binding Proteins from the Target Structure Alone. *Nature* **2022**, *605*, 551–560. [\[CrossRef\]](#) [\[PubMed\]](#)
- [58] Ingraham, J.B.; Baranov, M.; Costello, Z.; Barber, K.W.; Wang, W.; Ismail, A.; Frappier, V.; Lord, D.M.; Ng-Thow-Hing, C.; Van Vlack, E.R.; et al. Illuminating Protein Space with a Programmable Generative Model. *Nature* **2023**, *623*, 1070–1078. [\[CrossRef\]](#) [\[PubMed\]](#)
- [59] Ribeiro-Filho, H.V.; Jara, G.E.; Guerra, J.V.S.; Cheung, M.; Felbinger, N.R.; Pereira, J.G.C.; Pierce, B.G.; Lopes-De-Oliveira, P.S. Exploring the Potential of Structure-Based Deep Learning Approaches for T Cell Receptor Design. *PLoS Comput. Biol.* **2024**, *20*, e1012489. [\[CrossRef\]](#) [\[PubMed\]](#)
- [60] Meehl, M.M.; Immadisetty, K.; Trivedi, V.D.; Glowacki, P.; Prinzing, B.; Anido, A.A.; Ibañez-Vega, J.; Leslie, B.J.; Babu, M.M.; Krenciute, G. Computational Structural Optimization Enhances IL13R α 2—B7-H3 Tandem CAR T Cells to Overcome Antigen-Heterogeneity-Mediated Tumor Escape. *Mol. Ther.* **2025**, *33*, 4968–4987. [\[CrossRef\]](#)
- [61] Qiu, S.; Chen, J.; Wu, T.; Li, L.; Wang, G.; Wu, H.; Song, X.; Liu, X.; Wang, H. CAR-Toner: An AI-Driven Approach for CAR Tonic Signaling Prediction and Optimization. *Cell Res.* **2024**, *34*, 386–388. [\[CrossRef\]](#) [\[PubMed\]](#)
- [62] Nunes, F.V.M.; Behrens, L.M.P.; Weimer, R.D.; Gonçalves, G.F.; Fernandes, G.d.S.; Dorn, M. Deep Learning Methods and Applications in Single-Cell Multimodal Data Integration. *Mol. Omics* **2025**, *21*, 545–565. [\[CrossRef\]](#) [\[PubMed\]](#)
- [63] Liu, J.; Cen, X.; Yi, C.; Wang, F.-A.; Ding, J.; Cheng, J.; Wu, Q.; Gai, B.; Zhou, Y.; He, R.; et al. Challenges in AI-Driven Biomedical Multimodal Data Fusion and Analysis. *Genom. Proteom. Bioinform.* **2025**, *23*, qzaf011. [\[CrossRef\]](#) [\[PubMed\]](#)
- [64] Wu, X.; Yang, X.; Dai, Y.; Zhao, Z.; Zhu, J.; Guo, H.; Yang, R. Single-Cell Sequencing to Multi-Omics: Technologies and Applications. *Biomark. Res.* **2024**, *12*, 110. [\[CrossRef\]](#)
- [65] Lamiabile, A.; Champetier, T.; Leonardi, F.; Cohen, E.; Sommer, P.; Hardy, D.; Argy, N.; Massougboji, A.; Del Nery, E.; Cottrell, G.; et al. Revealing Invisible Cell Phenotypes with Conditional Generative Modeling. *Nat. Commun.* **2023**, *14*, 6386. [\[CrossRef\]](#)
- [66] Agarwal, A.A.; Harrang, J.; Noble, D.; McGowan, K.L.; Lange, A.W.; Engelhart, E.; Lahman, M.C.; Adamo, J.; Yu, X.; Serang, O.; et al. AlphaBind, a Domain-Specific Model to Predict and Optimize Antibody–Antigen Binding Affinity. *mAbs* **2025**, *17*, 2534626. [\[CrossRef\]](#)
- [67] Karthikeyan, D.; Bennett, S.N.; Reynolds, A.G.; Vincent, B.G.; Rubinsteyn, A. Conditional Generation of Real Antigen-Specific T Cell Receptor Sequences. *Nat. Mach. Intell.* **2025**, *7*, 1494–1509. [\[CrossRef\]](#)
- [68] Lawton, M.L.; Inge, M.M.; Blum, B.C.; Smith-Mahoney, E.L.; Bolzan, D.; Lin, W.; McConney, C.; Porter, J.; Moore, J.; Youssef, A.; et al. Mul-

- tiomic Profiling of Chronically Activated CD4⁺ T Cells Identifies Drivers of Exhaustion and Metabolic Reprogramming. *PLoS Biol.* **2024**, *22*, e3002943. [CrossRef]
- [69] Iu, D.S.; Maya, J.; Vu, L.T.; Fogarty, E.A.; McNairn, A.J.; Ahmed, F.; Franconi, C.J.; Munn, P.R.; Grenier, J.K.; Hanson, M.R.; et al. Transcriptional Reprogramming Primes CD8⁺ T Cells toward Exhaustion in Myalgic Encephalomyelitis/Chronic Fatigue Syndrome. *Proc. Natl. Acad. Sci. USA* **2024**, *121*, e2415119121. [CrossRef]
- [70] Hudson, D.; Fernandes, R.A.; Basham, M.; Ogg, G.; Koohy, H. Can We Predict T Cell Specificity with Digital Biology and Machine Learning? *Nat. Rev. Immunol.* **2023**, *23*, 511–521. [CrossRef]
- [71] Leggieri, P.A.; Liu, Y.; Hayes, M.; Connors, B.; Sepälä, S.; O'Malley, M.A.; Venturelli, O.S. Integrating Systems and Synthetic Biology to Understand and Engineer Microbiomes. *Annu. Rev. Biomed. Eng.* **2021**, *23*, 169–201. [CrossRef]
- [72] Leipzig, J.; Nüst, D.; Hoyt, C.T.; Ram, K.; Greenberg, J. The Role of Metadata in Reproducible Computational Research. *Patterns* **2021**, *2*, 100322. [CrossRef]
- [73] Tule, S.; Foley, G.; Bodén, M. Do Protein Language Models Learn Phylogeny? *Brief. Bioinform.* **2024**, *26*, bbaf047. [CrossRef]
- [74] Vieira, M.C.; Palm, A.-K.E.; Stamper, C.T.; Tepora, M.E.; Nguyen, K.D.; Pham, T.D.; Boyd, S.D.; Wilson, P.C.; Cobey, S. Germline-Encoded Specificities and the Predictability of the B Cell Response. *PLoS Pathog.* **2023**, *19*, e1011603. [CrossRef]
- [75] Javdan, S.B.; Deans, T.L. Design and Development of Engineered Receptors for Cell and Tissue Engineering. *Curr. Opin. Syst. Biol.* **2021**, *28*, 100363. [CrossRef]
- [76] Wong, W.W.; Lim, W.A. Golden Age of Immunoen지니어ing. *Immunol. Rev.* **2023**, *320*, 4–9. [CrossRef]
- [77] Rees, A.R. Understanding the Human Antibody Repertoire. *mAbs* **2020**, *12*, 1729683. [CrossRef]
- [78] Villanueva-Flores, F.; Sanchez-Villamil, J.I.; Garcia-Atutxa, I. Publisher Correction: AI-Driven Epitope Prediction: A Systematic Review, Comparative Analysis, and Practical Guide for Vaccine Development. *npj Vaccines* **2025**, *10*, 209. [CrossRef] [PubMed]
- [79] Elfatimi, E.; Lekbach, Y.; Prakash, S.; Ben-Mohamed, L. Artificial Intelligence and Machine Learning in the Development of Vaccines and Immunotherapeutics—Yesterday, Today, and Tomorrow. *Front. Artif. Intell.* **2025**, *8*, 1620572. [CrossRef] [PubMed]
- [80] Shugay, M.; Bagaev, D.V.; Zvyagin, I.V.; Vroomans, R.M.; Crawford, J.C.; Dolton, G.; Komech, E.A.; Sycheva, A.L.; Koneva, A.E.; Egorov, E.S.; et al. VDJdb: A Curated Database of T-Cell Receptor Sequences with Known Antigen Specificity. *Nucleic Acids Res.* **2018**, *46*, 419–427. [CrossRef]
- [81] Zhang, W.; Wang, L.; Liu, K.; Wei, X.; Yang, K.; Du, W.; Wang, S.; Guo, N.; Ma, C.; Luo, L.; et al. PIRD: Pan Immune Repertoire Database. *Bioinformatics* **2020**, *36*, 897–903. [CrossRef]
- [82] Textor, J.; Buytenhuijs, F.; Rogers, D.; Gauthier, È.M.; Sultan, S.; Wortel, I.M.; Kalies, K.; Fährnich, A.; Pagel, R.; Melichar, H.J.; et al. Machine Learning Analysis of the T Cell Receptor Repertoire Identifies Sequence Features of Self-Reactivity. *Cell Syst.* **2023**, *14*, 1059–1073. [CrossRef]
- [83] Greenshields-Watson, A.; Abanades, B.; Deane, C.M. Investigating the Ability of Deep Learning-Based Structure Prediction to Extrapolate and/or Enrich the Set of Antibody CDR Canonical Forms. *Front. Immunol.* **2024**, *15*, 1352703. [CrossRef] [PubMed]
- [84] Ostrovsky-Berman, M.; Frankel, B.; Polak, P.; Yaari, G. Immune2vec: Embedding B/T Cell Receptor Sequences in $\mathbb{R}^N$ Using Natural Language Processing. *Front. Immunol.* **2021**, *12*, 680687. [CrossRef] [PubMed]
- [85] Sidhom, J.-W.; Larman, H.B.; Pardoll, D.M.; Baras, A.S. DeepTCR Is a Deep Learning Framework for Revealing Sequence Concepts within T-Cell Repertoires. *Nat. Commun.* **2021**, *12*, 1605. [CrossRef]
- [86] Li, L.; Peng, X.; Batliwala, M.; Bouvier, M. Crystal Structures of MHC Class I Complexes Reveal the Elusive Intermediate Conformations Explored during Peptide Editing. *Nat. Commun.* **2023**, *14*, 5020. [CrossRef]
- [87] Sumida, K.H.; Núñez-Franco, R.; Kalvet, I.; Pellock, S.J.; Wicky, B.I.M.; Milles, L.F.; Dauparas, J.; Wang, J.; Kipnis, Y.; Jameson, N.; et al. Improving Protein Expression, Stability, and Function with Protein-MPNN. *J. Am. Chem. Soc.* **2024**, *146*, 2054–2061. [CrossRef]
- [88] Meng, F.; Zhou, N.; Hu, G.; Liu, R.; Zhang, Y.; Jing, M.; Hou, Q. A Comprehensive Overview of Recent Advances in Generative Models for Antibodies. *Comput. Struct. Biotechnol. J.* **2024**, *23*, 2648–2660. [CrossRef] [PubMed]
- [89] Lin, V.; Cheung, M.; Gowthaman, R.; Eisenberg, M.; Baker, B.M.; Pierce, B.G. TCR3d 2.0: Expanding the T Cell Receptor Structure Database with New Structures, Tools and Interactions. *Nucleic Acids Res.* **2025**, *53*, 604–608. [CrossRef]
- [90] Schmirler, R.; Heinzinger, M.; Rost, B. Fine-Tuning Protein Language Models Boosts Predictions across Diverse Tasks. *Nat. Commun.* **2024**, *15*, 7407. [CrossRef]
- [91] Drost, F.; An, Y.; Bonafonte-Pardàs, I.; Dratva, L.M.; Lindeboom, R.G.H.; Haniffa, M.; Teichmann, S.A.; Theis, F.; Lotfollahi, M.; Schubert, B. Multi-Modal Generative Modeling for Joint Analysis of Single-Cell T Cell Receptor and Gene Expression Data. *Nat. Commun.* **2024**, *15*, 5577. [CrossRef]
- [92] Abramson, J.; Adler, J.; Dunger, J.; Evans, R.; Green, T.; Pritzel, A.; Ronneberger, O.; Willmore, L.; Ballard, A.J.; Bambrick, J.; et al. Accurate Structure Pre-

- diction of Biomolecular Interactions with AlphaFold 3. *Nature* **2024**, *630*, 493–500. [[CrossRef](#)]
- [93] Deleuran, S.N.; Nielsen, M. NetTCR-struct, a Structure Driven Approach for Prediction of TCR-pMHC Interactions. *Front. Immunol.* **2025**, *16*, 1616328. [[CrossRef](#)]
- [94] Lourenço, A.; Subramanian, A.; Spencer, R.; Anaya, M.; Miao, J.; Fu, W.; Chow, E.; Thomson, M. Protein CREATE Enables Closed-Loop Design of De Novo Synthetic Protein Binders. *bioRxiv* **2025**. [[CrossRef](#)]
- [95] Yang, J.; Lal, R.G.; Bowden, J.C.; Astudillo, R.; Hameedi, M.A.; Kaur, S.; Hill, M.; Yue, Y.; Arnold, F.H. Active Learning-Assisted Directed Evolution. *Nat. Commun.* **2025**, *16*, 714. [[CrossRef](#)]
- [96] Gao, Y.; Gao, Y.; Wu, S.; Li, D.; Zhou, C.; Meng, F.; Dong, K.; Zhao, X.; Li, P.; Liang, A.; et al. Weakly Supervised Peptide-TCR Binding Prediction Facilitates Neoantigen Identification. *Cell Syst.* **2025**, *16*, 101403. [[CrossRef](#)] [[PubMed](#)]
- [97] Pertseva, M.; Follonier, O.; Scarcella, D.; Reddy, S.T. TCR Clustering by Contrastive Learning on Antigen Specificity. *Brief. Bioinform.* **2024**, *25*, bbae375. [[CrossRef](#)]
- [98] Saadat, M.; Zare-Mirakabad, F.; Masoudi-Nejad, A.; Baradaran, M.F.; Hosseinkhan, N. HLA-PepBinder: An Ensemble Model for the Prediction of HLA-Peptide Binding. *Iran. J. Biotechnol.* **2024**, *22*, e3927. [[CrossRef](#)]
- [99] Gioia, D.; Bertazzo, M.; Recanatini, M.; Masetti, M.; Cavalli, A. Dynamic Docking: A Paradigm Shift in Computational Drug Discovery. *Molecules* **2017**, *22*, 2029. [[CrossRef](#)]
- [100] Yilmaz, Ö.Z.G.; Doruker, P.; Kurkcuoglu, O. A Computationally Efficient Method to Generate Plausible Conformers for Ensemble Docking and Binding Free Energy Calculations. *J. Chem. Inf. Model.* **2025**, *65*, 8137–8157. [[CrossRef](#)] [[PubMed](#)]
- [101] Bielska, W.; Jaszczyszyn, I.; Dudzic, P.; Janusz, B.; Chomicz, D.; Wrobel, S.; Greiff, V.; Feehan, R.; Adolf-Bryfogle, J.; Krawczyk, K. Applying Computational Protein Design to Therapeutic Antibody Discovery—Current State and Perspectives. *Front. Immunol.* **2025**, *16*, 1571371. [[CrossRef](#)] [[PubMed](#)]
- [102] Sverchkov, Y.; Craven, M. A Review of Active Learning Approaches to Experimental Design for Uncovering Biological Networks. *PLoS Comput. Biol.* **2017**, *13*, e1005466. [[CrossRef](#)]
- [103] Slavny, P.; Hegde, M.; Doerner, A.; Parthiban, K.; McCafferty, J.; Zielonka, S.; Hoet, R. Advancements in Mammalian Display Technology for Therapeutic Antibody Development and Beyond: Current Landscape, Challenges, and Future Prospects. *Front. Immunol.* **2024**, *15*, 1469329. [[CrossRef](#)]
- [104] Ching, T.; Himmelstein, D.S.; Beaulieu-Jones, B.K.; Kalinin, A.A.; Do, B.T.; Way, G.P.; Ferrero, E.; Agapow, P.-M.; Zietz, M.; Hoffman, M.M.; et al. Opportunities and Obstacles for Deep Learning in Biology and Medicine. *J. R. Soc. Interface* **2018**, *15*, 20170387. [[CrossRef](#)]
- [105] Wossnig, L.; Furtmann, N.; Buchanan, A.; Kumar, S.; Greiff, V. Best Practices for Machine Learning in Antibody Discovery and Development. *Drug Discov. Today* **2024**, *29*, 104025. [[CrossRef](#)] [[PubMed](#)]
- [106] Rajagopal, N.; Choudhary, U.; Tsang, K.; Martin, K.P.; Karadag, M.; Chen, H.-T.; Kwon, N.-Y.; Mozdierz, J.; Horspool, A.M.; Li, L.; et al. Deep Learning-Based Design and Experimental Validation of a Medicine-Like Human Antibody Library. *Brief. Bioinform.* **2024**, *26*, bba023. [[CrossRef](#)]
- [107] Dixon, T.; MacPherson, D.; Mostofian, B.; Dazhenka, T.; Lotz, S.; McGee, D.; Shechter, S.; Shrestha, U.R.; Wiewiora, R.; McDargh, Z.A.; et al. Predicting the Structural Basis of Targeted Protein Degradation by Integrating Molecular Dynamics Simulations with Structural Mass Spectrometry. *Nat. Commun.* **2022**, *13*, 5884. [[CrossRef](#)]
- [108] Hati, S.; Bhattacharyya, S. Incorporating Modeling and Simulations in Undergraduate Biophysical Chemistry Course to Promote Understanding of Structure-Dynamics-Function Relationships in Proteins. *Biochem. Mol. Biol. Educ.* **2016**, *44*, 140–159. [[CrossRef](#)]
- [109] Hamamsy, T.; Barot, M.; Morton, J.T.; Steinegger, M.; Bonneau, R.; Cho, K. Learning Sequence, Structure, and Function Representations of Proteins with Language Models. *bioRxiv* **2023**. [[CrossRef](#)] [[PubMed](#)]
- [110] Dibaeinia, P.; Ojha, A.; Sinha, S. Interpretable AI for Inference of Causal Molecular Relationships from Omics Data. *Sci. Adv.* **2025**, *11*, eadk0837. [[CrossRef](#)]
- [111] Davidsen, K.; Olson, B.J.; DeWitt, W.S., 3rd; Feng, J.; Harkins, E.; Bradley, P.; Matsen, F.A., IV. Deep Generative Models for T Cell Receptor Protein Sequences. *eLife* **2019**, *8*, e46935. [[CrossRef](#)]
- [112] Isacchini, G.; Sethna, Z.; Elhanati, Y.; Nourmohammad, A.; Walczak, A.M.; Mora, T. Generative Models of T-Cell Receptor Sequences. *Phys. Rev. E* **2020**, *101*, 062414. [[CrossRef](#)]
- [113] Al-Jumaily, A.; Mukaidaisi, M.; Vu, A.; Tchagang, A.; Li, Y. Examining Multi-Objective Deep Reinforcement Learning Frameworks for Molecular Design. *Biosystems* **2023**, *232*, 104989. [[CrossRef](#)] [[PubMed](#)]
- [114] Wang, J.; Zhu, F. Multi-Objective Molecular Generation via Clustered Pareto-Based Reinforcement Learning. *Neural Netw.* **2024**, *179*, 106596. [[CrossRef](#)]
- [115] Niarakis, A.; Laubenbacher, R.; An, G.; Ilan, Y.; Fisher, J.; Flobak, Å.; Reiche, K.; Martínez, M.R.; Geris, L.; Ladeira, L.; et al. Immune Digital Twins for Complex Human Pathologies: Applications, Limitations, and Challenges. *npj Syst. Biol. Appl.* **2024**, *10*, 141. [[CrossRef](#)] [[PubMed](#)]
- [116] Zrimec, J.; Fu, X.; Muhammad, A.S.; Skrekas, C.; Jauniskis, V.; Speicher, N.K.; Börlin, C.S.; Verendel, V.; Chehreghani, M.H.; Dubhashi, D.; et al. Controlling Gene Expression with Deep Generative Design

- p>of Regulatory DNA.
- Nat. Commun.*
- 2022**
- ,
- 13*
- , 5099. [CrossRef] [PubMed]
- [117] Shahzadi, M.; Rafique, H.; Waheed, A.; Naz, H.; Waheed, A.; Zokirova, F.R.; Khan, H. Artificial Intelligence for Chimeric Antigen Receptor-Based Therapies: A Comprehensive Review of Current Applications and Future Perspectives. *Ther. Adv. Vaccines Immunother.* **2024**, *12*, 25151355241305856. [CrossRef]
- [118] Croce, G.; Lani, R.; Tardivon, D.; Bobisse, S.; de Tiani, M.; Bragina, M.; Perez, M.A.S.; Michaux, J.; Pak, H.S.; Michel, A.; et al. Phage Display Enables Machine Learning Discovery of Cancer Antigen-Specific TCRs. *Sci. Adv.* **2025**, *11*, eads5589. [CrossRef]
- [119] Fast, E.; Dhar, M.; Chen, B. TAPIR: A T-Cell Receptor Language Model for Predicting Rare and Novel Targets. *bioRxiv* **2023**. [CrossRef]
- [120] Semilietof, A.; Stefanidis, E.; Gray-Gaillard, E.; Pujol, J.; D'Esposito, A.; Reichenbach, P.; Guillaume, P.; Zoete, V.; Irving, M.; Michielin, O. Preclinical Model for Evaluating Human TCRs against Chimeric Syngeneic Tumors. *J. Immunother. Cancer* **2024**, *12*, e009504. [CrossRef]
- [121] Liu, B.; Greenwood, N.F.; Bonzanini, J.E.; Motmaen, A.; Meyerberg, J.; Dao, T.; Xiang, X.; Ault, R.; Sharp, J.; Wang, C.; et al. Design of High-Specificity Binders for Peptide-MHC-I Complexes. *Science* **2025**, *389*, 386–391. [CrossRef]
- [122] Johansen, K.H.; Wolff, D.S.; Scapolo, B.; Fernández-Quintero, M.L.; Christensen, C.R.; Loeffler, J.R.; Rivera-de-Torre, E.; Overath, M.D.; Munk, K.K.; Morell, O.; et al. De Novo-Designed pMHC Binders Facilitate T Cell-Mediated Cytotoxicity toward Cancer Cells. *Science* **2025**, *389*, 380–385. [CrossRef]
- [123] Linette, G.P.; Stadtmauer, E.A.; Maus, M.V.; Rapoport, A.P.; Levine, B.L.; Emery, L.; Litzky, L.; Bagg, A.; Carreno, B.M.; Cimino, P.J.; et al. Cardiovascular Toxicity and Titin Cross-Reactivity of Affinity-Enhanced T Cells in Myeloma and Melanoma. *Blood* **2013**, *122*, 863–871. [CrossRef]
- [124] Cameron, B.J.; Gerry, A.B.; Dukes, J.; Harper, J.V.; Kannan, V.; Bianchi, F.C.; Grand, F.; Brewer, J.E.; Gupta, M.; Plesa, G.; et al. Identification of a Titin-Derived HLA-A1–Presented Peptide as a Cross-Reactive Target for Engineered MAGE A3–Directed T Cells. *Sci. Transl. Med.* **2013**, *5*, 197ra103. [CrossRef]
- [125] Foldvari, Z.; Knetter, C.; Yang, W.; Gjerdingen, T.J.; Bollineni, R.C.; Tran, T.T.; Lund-Johansen, F.; Kolstad, A.; Drosch, K.; Klopffleisch, R.; et al. A Systematic Safety Pipeline for Selection of T-Cell Receptors to Enter Clinical Use. *npj Vaccines* **2023**, *8*, 126. [CrossRef]
- [126] Reynisson, B.; Alvarez, B.; Paul, S.; Peters, B.; Nielsen, M. 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.* **2020**, *48*, 449–454. [CrossRef]
- [127] Li, G.; Iyer, B.; Prasath, V.B.S.; Ni, Y.; Salomonis, N. DeepImmuno: Deep Learning-Empowered Prediction and Generation of Immunogenic Peptides for T-Cell Immunity. *Brief. Bioinform.* **2021**, *22*, bbab160. [CrossRef]
- [128] Bradley, P. Structure-Based Prediction of T Cell Receptor:Peptide-MHC Interactions. *eLife* **2023**, *12*, e82813. [CrossRef]
- [129] Ishii, K.; Davies, J.S.; Sinkoe, A.L.; Nguyen, K.A.; Norberg, S.M.; McIntosh, C.P.; Kadakia, T.; Serna, C.; Rae, Z.; Kelly, M.C.; et al. Multi-Tiered Approach to Detect Autoimmune Cross-Reactivity of Therapeutic T Cell Receptors. *Sci. Adv.* **2023**, *9*, eadg9845. [CrossRef]
- [130] Rollins, Z.A.; Faller, R.; George, S.C. Using Molecular Dynamics Simulations to Interrogate T Cell Receptor Non-Equilibrium Kinetics. *Comput. Struct. Biotechnol. J.* **2022**, *20*, 2124–2133. [CrossRef]
- [131] Mazhibiyeva, A.; Pham, T.T.; Pats, K.; Lukac, M.; Molnár, F. Bridging Prediction and Reality: Comprehensive Analysis of Experimental and AlphaFold 2 Full-Length Nuclear Receptor Structures. *Comput. Struct. Biotechnol. J.* **2025**, *27*, 1998–2013. [CrossRef]
- [132] Yin, R.; Ribeiro-Filho, H.V.; Lin, V.; Gowthaman, R.; Cheung, M.; Pierce, B.G. TCRmodel2: High-Resolution Modeling of T Cell Receptor Recognition Using Deep Learning. *Nucleic Acids Res.* **2023**, *51*, 569–576. [CrossRef] [PubMed]
- [133] Krishna, R.; Wang, J.; Ahern, W.; Sturmfels, P.; Venkatesh, P.; Kalvet, I.; Lee, G.R.; Morey-Burrows, F.S.; Anishchenko, I.; Humphreys, I.R.; et al. Generalized Biomolecular Modeling and Design with RoseTTAFold All-Atom. *Science* **2024**, *384*, eadl2528. [CrossRef] [PubMed]
- [134] Robertson, I.B.; Mulvaney, R.; Dieckmann, N.; Vantellini, A.; Canestraro, M.; Amicarella, F.; O'Dwyer, R.; Cole, D.K.; Harper, S.; Dushek, O.; et al. Tuning the Potency and Selectivity of ImmTAC Molecules by Affinity Modulation. *Clin. Exp. Immunol.* **2024**, *215*, 105–119. [CrossRef]
- [135] Sinclair, A.; Krämer, S.; Reinhart, C.; Stehle, J.; Schuster, S.; Herz, T.; Al Hasani, H.; Hamde, P.; Selinger, O.; Birkenfeld, J. Beyond Sequence Similarity: ML-Powered Identification of pHLA Off-Targets for TCR-Mimic Antibodies Using High Throughput Binding Kinetics. *mAbs* **2026**, *18*, 2601360. [CrossRef] [PubMed]
- [136] Wellach, K.; Riemer, A.B. Highly Sensitive Live-Cell Imaging-Based Cytotoxicity Assay Enables Functional Validation of Rare Epitope-Specific CTLs. *Front. Immunol.* **2025**, *16*, 1558620. [CrossRef]
- [137] Foulke, J.G.; Chen, L.; Chang, H.; McManus, C.E.; Tian, F.; Gu, Z. Optimizing Ex Vivo CAR-T Cell-Mediated Cytotoxicity Assay through Multimodality Imaging. *Cancers* **2024**, *16*, 2497. [CrossRef]
- [138] Banerjee, A.; Pattinson, D.J.; Wincek, C.L.; Bunk, P.; Axhemi, A.; Chapin, S.R.; Navlakha, S.; Meyer,

- H.V. T Cell Receptor Cross-Reactivity Prediction Improved by a Comprehensive Mutational Scan Database. *Cell Syst.* **2025**, *16*, 101345. [[CrossRef](#)]
- [139] Sharma, G.; Teng, F.; Round, J.; Sneddon, S.A.; Sivasothy, S.; Brown, S.D.; Holt, R.A. Comprehensive Self-Antigen Screening to Assess Cross-Reactivity in Promiscuous T-Cell Receptors. *Front. Immunol.* **2025**, *16*, 1719827. [[CrossRef](#)]
- [140] Tang, X.; Dai, H.; Knight, E.; Wu, F.; Li, Y.; Li, T.; Gerstein, M. A Survey of Generative AI for De Novo Drug Design: New Frontiers in Molecule and Protein Generation. *Brief. Bioinform.* **2024**, *25*, bbae338. [[CrossRef](#)]
- [141] Borrmann, T.; Pierce, B.G.; Vreven, T.; Baker, B.M.; Weng, Z. High-Throughput Modeling and Scoring of TCR-pMHC Complexes to Predict Cross-Reactive Peptides. *Bioinformatics* **2021**, *36*, 5377–5385. [[CrossRef](#)]
- [142] Seo, S.Y.; Rhee, J.-K. TCR-epiDiff: Solving Dual Challenges of TCR Generation and Binding Prediction. *Bioinformatics* **2025**, *41*, 125–132. [[CrossRef](#)]
- [143] Kalematis, M.; Noroozi, A.; Shahbakhsh, A.; Koohi, S. ParaAntiProt Provides Paratope Prediction Using Antibody and Protein Language Models. *Sci. Rep.* **2024**, *14*, 29141. [[CrossRef](#)]
- [144] Ghanbargpour, A.; Jiang, M.; Foster, D.; Chai, Q. Structure-Free Antibody Paratope Similarity Prediction for in Silico Epitope Binning via Protein Language Models. *iScience* **2023**, *26*, 106036. [[CrossRef](#)]
- [145] Pruvost, T.; Mathieu, M.; Dubois, S.; Maillère, B.; Vigne, E.; Nozach, H. Deciphering Cross-Species Reactivity of LAMP-1 Antibodies Using Deep Mutational Epitope Mapping and AlphaFold. *mAbs* **2023**, *15*, 2175311. [[CrossRef](#)]
- [146] Mason, D.M.; Reddy, S.T. Predicting Adaptive Immune Receptor Specificities by Machine Learning Is a Data Generation Problem. *Cell Syst.* **2024**, *15*, 1190–1197. [[CrossRef](#)] [[PubMed](#)]
- [147] Abeer, A.N.M.N.; Urban, N.M.; Weil, M.R.; Alexander, F.J.; Yoon, B.-J. Multi-Objective Latent Space Optimization of Generative Molecular Design Models. *Patterns* **2024**, *5*, 101042. [[CrossRef](#)] [[PubMed](#)]
- [148] Lindner, S.E.; Johnson, S.M.; Brown, C.E.; Wang, L.D. Chimeric Antigen Receptor Signaling: Functional Consequences and Design Implications. *Sci. Adv.* **2020**, *6*, eaaz3223. [[CrossRef](#)]
- [149] Guedan, S.; Calderon, H.; Posey, A.D., Jr.; Maus, M.V. Engineering and Design of Chimeric Antigen Receptors. *Mol. Ther. Methods Clin. Dev.* **2019**, *12*, 145–156. [[CrossRef](#)]
- [150] Castellanos-Rueda, R.; Wang, K.-L.K.; Forster, J.L.; Driessen, A.; Frank, J.A.; Martínez, M.R.; Reddy, S.T. Dissecting the Role of CAR Signaling Architectures on T Cell Activation and Persistence Using Pooled Screens and Single-Cell Sequencing. *Sci. Adv.* **2025**, *11*, eadp4008. [[CrossRef](#)] [[PubMed](#)]
- [151] Smirnov, S.; Mateikovich, P.; Samochernykh, K.; Shlyakhto, E. Recent Advances on CAR-T Signaling Pave the Way for Prolonged Persistence and New Modalities in Clinic. *Front. Immunol.* **2024**, *15*, 1335424. [[CrossRef](#)]
- [152] Alsaieedi, A.A.; Zaher, K.A. Tracing the Development of CAR-T Cell Design: From Concept to Next-Generation Platforms. *Front. Immunol.* **2025**, *16*, 1615212. [[CrossRef](#)]
- [153] Sutanto, H.; Fetarayani, D. Integrating Artificial Intelligence into Small Molecule Development for Precision Cancer Immunomodulation Therapy. *npj Drug Discov.* **2025**, *2*, 25. [[CrossRef](#)]
- [154] Du, Q.; Wang, H.; Jiang, B.; Wang, X. Advancing Genetic Engineering with Active Learning: Theory, Implementations and Potential Opportunities. *Brief. Bioinform.* **2025**, *26*, bbaf286. [[CrossRef](#)]
- [155] Ferdous, S.; Shihab, I.F.; Chowdhury, R.; Reuel, N.F. Reinforcement Learning-Guided Control Strategies for CAR T-Cell Activation and Expansion. *Biotechnol. Bioeng.* **2024**, *121*, 2868–2880. [[CrossRef](#)] [[PubMed](#)]
- [156] Chandrasekar, V.; Panicker, A.J.; Dey, A.K.; Mohammad, S.; Chakraborty, A.; Samal, S.K.; Dash, A.; Bhadra, J.; Suar, M.; Khare, M.; et al. Integrated Approaches for Immunotoxicity Risk Assessment: Challenges and Future Directions. *Discov. Toxicol.* **2024**, *1*, 9. [[CrossRef](#)]
- [157] Yeku, O.O.; Brentjens, R.J. Armored CAR T-Cells: Utilizing Cytokines and Pro-Inflammatory Ligands to Enhance CAR T-Cell Anti-Tumour Efficacy. *Biochem. Soc. Trans.* **2016**, *44*, 412–418. [[CrossRef](#)] [[PubMed](#)]
- [158] Li, X.; Chen, T.; Li, X.; Zhang, H.; Li, Y.; Zhang, S.; Luo, S.; Zheng, T. Therapeutic Targets of Armored Chimeric Antigen Receptor T Cells Navigating the Tumor Microenvironment. *Exp. Hematol. Oncol.* **2024**, *13*, 96. [[CrossRef](#)]
- [159] Han, X.; Wang, Y.; Wei, J.; Han, W. Multi-Antigen-Targeted Chimeric Antigen Receptor T Cells for Cancer Therapy. *J. Hematol. Oncol.* **2019**, *12*, 128. [[CrossRef](#)]
- [160] Singh, A.V.; Bhardwaj, P.; Laux, P.; Pradeep, P.; Busse, M.; Luch, A.; Hirose, A.; Osgood, C.J.; Stacey, M.W. AI and ML-Based Risk Assessment of Chemicals: Predicting Carcinogenic Risk from Chemical-Induced Genomic Instability. *Front. Toxicol.* **2024**, *6*, 1461587. [[CrossRef](#)]
- [161] Zitnik, M.; Li, M.M.; Wells, A.; Glass, K.; Gysi, D.M.; Krishnan, A.; Murali, T.M.; Radivojac, P.; Roy, S.; Baudot, A.; et al. Current and Future Directions in Network Biology. *Bioinform. Adv.* **2024**, *4*, vbae099. [[CrossRef](#)]
- [162] Blay, V.; Pandiella, A. Strategies to Boost Antibody Selectivity in Oncology. *Trends Pharmacol. Sci.* **2024**, *45*, 1135–1149. [[CrossRef](#)]
- [163] This, S.; Costantino, S.; Melichar, H.J. Machine Learning Predictions of T Cell Antigen Specificity from Intracellular Calcium Dynamics. *Sci. Adv.* **2024**, *10*, eadk2298. [[CrossRef](#)] [[PubMed](#)]

- [164] Putignano, G.; Ruipérez-Campillo, S.; Yuan, Z.; Millet, J.; Guerrero-Aspizua, S. Mathematical Models and Computational Approaches in CAR-T Therapeutics. *Front. Immunol.* **2025**, *16*, 1581210. [[CrossRef](#)] [[PubMed](#)]
- [165] McFaline-Figueroa, J.L.; Srivatsan, S.; Hill, A.J.; Gasperini, M.; Jackson, D.L.; Saunders, L.; Domcke, S.; Regalado, S.G.; Lazarchuck, P.; Alvarez, S.; et al. Multiplex Single-Cell Chemical Genomics Reveals the Kinase Dependence of the Response to Targeted Therapy. *Cell Genom.* **2024**, *4*, 100487. [[CrossRef](#)]
- [166] Biederstädt, A.; Manzar, G.S.; Daher, M. Multiplexed Engineering and Precision Gene Editing in Cellular Immunotherapy. *Front. Immunol.* **2022**, *13*, 1063303. [[CrossRef](#)]
- [167] Zambrano, N.; Froehlich, G.; Lazarevic, D.; Pasariello, M.; Nicosia, A.; De Lorenzo, C.; Morelli, M.J.; Sasso, E. High-Throughput Monoclonal Antibody Discovery from Phage Libraries: Challenging the Current Preclinical Pipeline to Keep the Pace with the Increasing mAb Demand. *Cancers* **2022**, *14*, 1325. [[CrossRef](#)] [[PubMed](#)]
- [168] Lopez, R.; Regier, J.; Cole, M.B.; Jordan, M.I.; Yosef, N. Deep Generative Modeling for Single-Cell Transcriptomics. *Nat. Methods* **2018**, *15*, 1053–1058. [[CrossRef](#)]
- [169] Specht, H.; Emmott, E.; Petelski, A.A.; Huffman, R.G.; Perlman, D.H.; Serra, M.; Kharchenko, P.; Koller, A.; Slavov, N. Single-Cell Proteomic and Transcriptomic Analysis of Macrophage Heterogeneity Using SCope2. *Genome Biol.* **2021**, *22*, 50. [[CrossRef](#)]
- [170] Vijayan, R.; Kihlberg, J.; Cross, J.B.; Poongavanam, V. Enhancing Preclinical Drug Discovery with Artificial Intelligence. *Drug Discov. Today* **2022**, *27*, 967–984. [[CrossRef](#)]
- [171] Audagnotto, M.; Czechtizky, W.; De Maria, L.; Käck, H.; Papoian, G.; Tornberg, L.; Tyrchan, C.; Ulander, J. Machine Learning/Molecular Dynamic Protein Structure Prediction Approach to Investigate the Protein Conformational Ensemble. *Sci. Rep.* **2022**, *12*, 10018. [[CrossRef](#)]
- [172] Chen, L.; Balabanidou, V.; Remeta, D.P.; Minetti, C.A.; Portaliou, A.G.; Economou, A.; Kalodimos, C.G. Structural Instability Tuning as a Regulatory Mechanism in Protein-Protein Interactions. *Mol. Cell* **2011**, *44*, 734–744. [[CrossRef](#)] [[PubMed](#)]
- [173] Varadi, M.; Anyango, S.; Deshpande, M.; Nair, S.; Natassia, C.; Yordanova, G.; Yuan, D.; Stroe, O.; Wood, G.; Laydon, A.; et al. AlphaFold Protein Structure Database: Massively Expanding the Structural Coverage of Protein-Sequence Space with High-Accuracy Models. *Nucleic Acids Res.* **2022**, *50*, D439–D444. [[CrossRef](#)]
- [174] Maguire, J.B.; Haddox, H.K.; Strickland, D.; Halabiya, S.F.; Coventry, B.; Griffin, J.R.; Pulavarti, S.V.S.R.K.; Cummins, M.; Thieker, D.F.; Klavins, E.; et al. Perturbing the Energy Landscape for Improved Packing during Computational Protein Design. *Proteins Struct. Funct. Bioinform.* **2021**, *89*, 436–449. [[CrossRef](#)] [[PubMed](#)]
- [175] Eastman, P.; Swails, J.; Chodera, J.D.; McGibbon, R.T.; Zhao, Y.; Beauchamp, K.A.; Wang, L.-P.; Simonett, A.C.; Harrigan, M.P.; Stern, C.D.; et al. OpenMM 7: Rapid Development of High Performance Algorithms for Molecular Dynamics. *PLoS Comput. Biol.* **2017**, *13*, e1005659. [[CrossRef](#)] [[PubMed](#)]
- [176] Lyskov, S.; Gray, J.J. The RosettaDock Server for Local Protein-Protein Docking. *Nucleic Acids Res.* **2008**, *36*, 233–238. [[CrossRef](#)]
- [177] Pratt, O.S.; Elliott, L.G.; Haon, M.; Mesdaghi, S.; Price, R.M.; Simpkin, A.J.; Rigden, D.J. AlphaFold 2, but Not AlphaFold 3, Predicts Confident but Unrealistic β -Solenoid Structures for Repeat Proteins. *Comput. Struct. Biotechnol. J.* **2025**, *27*, 467–477. [[CrossRef](#)]
- [178] Chang, M.P.; Jin, T.; Gudinas, A.P.; Fernandez, D.; Alexander-Katz, A.; Matsui, T.; Mai, D.J. Repetitive Proteins That Undergo Large Conformational Changes Evade Structural Prediction Algorithms. *J. Chem. Phys.* **2025**, *163*, 224906. [[CrossRef](#)]
- [179] Dilchert, J.; Hofmann, M.; Unverdorben, F.; Kontermann, R.; Bunk, S. Mammalian Display Platform for the Maturation of Bispecific TCR-Based Molecules. *Antibodies* **2022**, *11*, 34. [[CrossRef](#)]
- [180] Smith, S.N.; Harris, D.T.; Kranz, D.M. T Cell Receptor Engineering and Analysis Using the Yeast Display Platform. *Methods Mol. Biol.* **2015**, *1319*, 95–141. [[CrossRef](#)]
- [181] Gai, S.A.; Wittrup, K.D. Yeast Surface Display for Protein Engineering and Characterization. *Curr. Opin. Struct. Biol.* **2007**, *17*, 467–473. [[CrossRef](#)]
- [182] Hartung, T. AI, Agentic Models and Lab Automation for Scientific Discovery—The Beginning of scAIInce. *Front. Artif. Intell.* **2025**, *8*, 1649155. [[CrossRef](#)]
- [183] Irvine, D.J.; Maus, M.V.; Mooney, D.J.; Wong, W.W. The Future of Engineered Immune Cell Therapies. *Science* **2022**, *378*, 853–858. [[CrossRef](#)]
- [184] Wang, W.; Hariharan, M.; Ding, W.; Bartlett, A.; Barragan, C.; Castanon, R.; Wang, R.; Rothenberg, V.; Song, H.; Nery, J.R.; et al. Genetics and Environment Distinctively Shape the Human Immune Cell Epigenome. *Nat. Genet.* **2025**, *58*, 392–403. [[CrossRef](#)] [[PubMed](#)]
- [185] Kondilis-Mangum, H.D.; Wade, P.A. Epigenetics and the Adaptive Immune Response. *Mol. Asp. Med.* **2013**, *34*, 813–825. [[CrossRef](#)] [[PubMed](#)]
- [186] Sadria, M.; Layton, A. scVAEDer: Integrating Deep Diffusion Models and Variational Autoencoders for Single-Cell Transcriptomics Analysis. *Genome Biol.* **2025**, *26*, 64. [[CrossRef](#)] [[PubMed](#)]
- [187] Jung, S. Advances in Modeling Cellular State Dynamics: Integrating Omics Data and Predictive Techniques. *Anim. Cells Syst.* **2025**, *29*, 72–83. [[CrossRef](#)]

- [188] Rodov, A.; Baniadam, H.; Zeiser, R.; Amit, I.; Yosef, N.; Wertheimer, T.; Ingelfinger, F. Towards the Next Generation of Data-Driven Therapeutics Using Spatially Resolved Single-Cell Technologies and Generative AI. *Eur. J. Immunol.* **2025**, *55*, e202451234. [\[CrossRef\]](#)
- [189] Ding, J.; Regev, A. Deep Generative Model Embedding of Single-Cell RNA-Seq Profiles on Hyperspheres and Hyperbolic Spaces. *Nat. Commun.* **2021**, *12*, 2554. [\[CrossRef\]](#)
- [190] Choi, H.; Kim, H.; Chung, H.; Lee, D.-S.; Kim, J. Application of Computational Algorithms for Single-Cell RNA-seq and ATAC-seq in Neurodegenerative Diseases. *Brief. Funct. Genom.* **2025**, *24*, elae044. [\[CrossRef\]](#)
- [191] Wang, C.; Liu, Z.-P. Diffusion-Based Generation of Gene Regulatory Networks from scRNA-seq Data with DigNet. *Genome Res.* **2025**, *35*, 340–354. [\[CrossRef\]](#)
- [192] Yeo, G.H.T.; Saksena, S.D.; Gifford, D.K. Generative Modeling of Single-Cell Time Series with PRESCIENT Enables Prediction of Cell Trajectories with Interventions. *Nat. Commun.* **2021**, *12*, 3222. [\[CrossRef\]](#)
- [193] Weerakoon, H.; Mohamed, A.; Wong, Y.; Chen, J.; Senadheera, B.; Haigh, O.; Watkins, T.S.; Kazakoff, S.; Mukhopadhyay, P.; Mulvenna, J.; et al. Integrative Temporal Multi-Omics Reveals Uncoupling of Transcriptome and Proteome during Human T Cell Activation. *npj Syst. Biol. Appl.* **2024**, *10*, 21. [\[CrossRef\]](#)
- [194] Wang, X.; Fan, D.; Yang, Y.; Gimple, R.C.; Zhou, S. Integrative Multi-Omics Approaches to Explore Immune Cell Functions: Challenges and Opportunities. *iScience* **2023**, *26*, 106359. [\[CrossRef\]](#)
- [195] Singh, A.V.; Bhardwaj, P.; Kishore, V.; Choudhary, S.; Hirose, A.; Gupta, N.; Busse, M.; Singh, S.L.; Osgood, C.J. Multifunctional Nanoscale Pigments: Emerging Risks and Circular Strategies for a Sustainable Future. *Small Sci.* **2025**, *5*, e202500240. [\[CrossRef\]](#) [\[PubMed\]](#)
- [196] Lotfollahi, M.; Susmelj, A.K.; De Donno, C.; Hetzel, L.; Ji, Y.; Ibarra, I.L.; Srivatsan, S.R.; Naghipourfar, M.; Daza, R.M.; Martin, B.; et al. Predicting Cellular Responses to Complex Perturbations in High-Throughput Screens. *Mol. Syst. Biol.* **2023**, *19*, e11517. [\[CrossRef\]](#)
- [197] Lotfollahi, M.; Wolf, F.A.; Theis, F.J. scGen Predicts Single-Cell Perturbation Responses. *Nat. Methods* **2019**, *16*, 715–721. [\[CrossRef\]](#) [\[PubMed\]](#)
- [198] Metzner, E.; Southard, K.M.; Norman, T.M. Multi-ome Perturb-seq Unlocks Scalable Discovery of Integrated Perturbation Effects on the Transcriptome and Epigenome. *Cell Syst.* **2025**, *16*, 101161. [\[CrossRef\]](#)
- [199] Zhou, P.; Shi, H.; Huang, H.; Sun, X.; Yuan, S.; Chapman, N.M.; Connelly, J.P.; Lim, S.A.; Saravia, J.; Kc, A.; et al. Single-Cell CRISPR Screens in Vivo Map T Cell Fate Regulomes in Cancer. *Nature* **2023**, *624*, 154–163. [\[CrossRef\]](#) [\[PubMed\]](#)
- [200] Belk, J.A.; Yao, W.; Ly, N.; Freitas, K.A.; Chen, Y.-T.; Shi, Q.; Valencia, A.M.; Shifrut, E.; Kale, N.; Yost, K.E.; et al. Genome-Wide CRISPR Screens of T Cell Exhaustion Identify Chromatin Remodeling Factors That Limit T Cell Persistence. *Cancer Cell* **2022**, *40*, 768–786. [\[CrossRef\]](#)
- [201] Xia, L.; Komissarova, A.; Jacover, A.; Shovman, Y.; Arcila-Barrera, S.; Tornovsky-Babeay, S.; Prakashan, M.M.J.; Nasereddin, A.; Plaschkes, I.; Nevo, Y.; et al. Systematic Identification of Gene Combinations to Target in Innate Immune Cells to Enhance T Cell Activation. *Nat. Commun.* **2023**, *14*, 6295. [\[CrossRef\]](#)
- [202] Prochazka, L.; Michaels, Y.S.; Lau, C.; Jones, R.D.; Siu, M.; Yin, T.; Wu, D.; Jang, E.; Vázquez-Cantú, M.; Gilbert, P.M.; et al. Synthetic Gene Circuits for Cell State Detection and Protein Tuning in Human Pluripotent Stem Cells. *Mol. Syst. Biol.* **2022**, *18*, e10886. [\[CrossRef\]](#) [\[PubMed\]](#)
- [203] Solé, R.; Conde-Pueyo, N.; Pla-Mauri, J.; García-Ojalvo, J.; Montserrat, N.; Levin, M. Open Problems in Synthetic Multicellularity. *npj Syst. Biol. Appl.* **2024**, *10*, 151. [\[CrossRef\]](#)
- [204] Nilsson, A.; Peters, J.M.; Meimetis, N.; Bryson, B.; Lauffenburger, D.A. Artificial Neural Networks Enable Genome-Scale Simulations of Intracellular Signaling. *Nat. Commun.* **2022**, *13*, 3069. [\[CrossRef\]](#)
- [205] Racovita, A.; Jaramillo, A. Reinforcement Learning in Synthetic Gene Circuits. *Biochem. Soc. Trans.* **2020**, *48*, 1637–1643. [\[CrossRef\]](#) [\[PubMed\]](#)
- [206] Ding, X.; Zhang, L.; Fan, M.; Li, L. Network-Based Transfer of Pan-Cancer Immunotherapy Responses to Guide Breast Cancer Prognosis. *npj Syst. Biol. Appl.* **2025**, *11*, 4. [\[CrossRef\]](#)
- [207] DaSilva, L.F.; Senan, S.; Patel, Z.M.; Reddy, A.J.; Gabbita, S.; Nussbaum, Z.; Córdova, C.M.V.; Wentele, A.; Weber, N.; Tunjic, T.M.; et al. DNA-Diffusion: Leveraging Generative Models for Controlling Chromatin Accessibility and Gene Expression via Synthetic Regulatory Elements. *bioRxiv* **2024**. [\[CrossRef\]](#)
- [208] Garcia, B.T.; Westerfield, L.; Yelemali, P.; Gogate, N.; Rivera-Munoz, E.A.; Du, H.; Dawood, M.; Jolly, A.; Lupski, J.R.; Posey, J.E. Improving Automated Deep Phenotyping through Large Language Models Using Retrieval-Augmented Generation. *Genome Med.* **2025**, *17*, 91. [\[CrossRef\]](#)
- [209] Dong, M.; Wang, L.; Hu, N.; Rao, Y.; Wang, Z.; Zhang, Y. Integration of Multi-Omics Approaches in Exploring Intra-Tumoral Heterogeneity. *Cancer Cell Int.* **2025**, *25*, 317. [\[CrossRef\]](#) [\[PubMed\]](#)
- [210] van Lent, P.; Schmitz, J.; Abeel, T. Simulated Design–Build–Test–Learn Cycles for Consistent Comparison of Machine Learning Methods in Metabolic Engineering. *ACS Synth. Biol.* **2023**, *12*, 2588–2599. [\[CrossRef\]](#)
- [211] Kitano, S.; Lin, C.; Foo, J.L.; Chang, M.W. Synthetic Biology: Learning the Way toward High-Precision Biological Design. *PLoS Biol.* **2023**, *21*, e3002116. [\[CrossRef\]](#)

- [212] Daniszewski, M.; Crombie, D.E.; Henderson, R.; Liang, H.H.; Wong, R.C.B.; Hewitt, A.W.; Pébay, A. Automated Cell Culture Systems and Their Applications to Human Pluripotent Stem Cell Studies. *JALA J. Assoc. Lab. Autom.* **2018**, *23*, 315–325. [[CrossRef](#)]
- [213] Lei, Y.; Tsang, J.S. Systems Human Immunology and AI: Immune Setpoint and Immune Health. *Annu. Rev. Immunol.* **2025**, *43*, 693–722. [[CrossRef](#)]
- [214] Gurdo, N.; Volke, D.C.; McCloskey, D.; Nikel, P.I. Automating the Design-Build-Test-Learn Cycle towards Next-Generation Bacterial Cell Factories. *New Biotechnol.* **2023**, *74*, 1–15. [[CrossRef](#)]
- [215] Si, Y.; Zou, J.; Gao, Y.; Chuai, G.; Liu, Q.; Chen, L. Foundation Models in Molecular Biology. *Biophys. Rep.* **2024**, *10*, 135–151. [[CrossRef](#)] [[PubMed](#)]
- [216] Moldwin, A.; Shehu, A. Foundation Models for AI-Enabled Biological Design. *arXiv* **2025**. [[CrossRef](#)]
- [217] Wytock, T.P.; Motter, A.E. Cell Reprogramming Design by Transfer Learning of Functional Transcriptional Networks. *Proc. Natl. Acad. Sci. USA* **2024**, *121*, e2312942121. [[CrossRef](#)] [[PubMed](#)]
- [218] Blaby, I.K.; Cheng, J.-F. Building a Custom High-Throughput Platform at the Joint Genome Institute for DNA Construct Design and Assembly—Present and Future Challenges. *Synth. Biol.* **2020**, *5*, ysaa023. [[CrossRef](#)]
- [219] Ma, Y.; Zhang, Z.; Jia, B.; Yuan, Y. Automated High-Throughput DNA Synthesis and Assembly. *Heliyon* **2024**, *10*, e26967. [[CrossRef](#)]
- [220] Rezalotfi, A.; Fritz, L.; Förster, R.; Bošnjak, B. Challenges of CRISPR-Based Gene Editing in Primary T Cells. *Int. J. Mol. Sci.* **2022**, *23*, 1689. [[CrossRef](#)] [[PubMed](#)]
- [221] Kwak, S.; Lee, H.; Yu, D.; Jeon, T.-J.; Kim, S.M.; Ryu, H. Microfluidic Platforms for Ex Vivo and In Vivo Gene Therapy. *Biosensors* **2025**, *15*, 504. [[CrossRef](#)]
- [222] Hunt, A.C.; Rasor, B.J.; Seki, K.; Ekas, H.M.; Warfel, K.F.; Karim, A.S.; Jewett, M.C. Cell-Free Gene Expression: Methods and Applications. *Chem. Rev.* **2025**, *125*, 91–149. [[CrossRef](#)]
- [223] Lu, Y. The Future of Cell-Free Synthetic Biology. *Biotechnol. Notes* **2024**, *5*, A1–A3. [[CrossRef](#)]
- [224] Craig, T.; Holland, R.; D’amore, R.; Johnson, J.R.; McCue, H.V.; West, A.; Zulkower, V.; Tekotte, H.; Cai, Y.; Swan, D.; et al. Leaf LIMS: A Flexible Laboratory Information Management System with a Synthetic Biology Focus. *ACS Synth. Biol.* **2017**, *6*, 2273–2280. [[CrossRef](#)]
- [225] Zhou, Y.; Shao, N.; de Castro, R.B.; Zhang, P.; Ma, Y.; Liu, X.; Huang, F.; Wang, R.-F.; Qin, L. Evaluation of Single-Cell Cytokine Secretion and Cell-Cell Interactions with a Hierarchical Loading Microwell Chip. *Cell Rep.* **2020**, *31*, 107574. [[CrossRef](#)]
- [226] Lim, J.; Park, C.; Kim, M.; Kim, H.; Kim, J.; Lee, D.-S. Advances in Single-Cell Omics and Multiomics for High-Resolution Molecular Profiling. *Exp. Mol. Med.* **2024**, *56*, 515–526. [[CrossRef](#)] [[PubMed](#)]
- [227] Roth, J.P.; Bajorath, J. Relationship between Prediction Accuracy and Uncertainty in Compound Potency Prediction Using Deep Neural Networks and Control Models. *Sci. Rep.* **2024**, *14*, 6536. [[CrossRef](#)] [[PubMed](#)]
- [228] Hillson, N.; Caddick, M.; Cai, Y.; Carrasco, J.A.; Chang, M.W.; Curach, N.C.; Bell, D.J.; Le Feuvre, R.; Friedman, D.C.; Fu, X.; et al. Building a Global Alliance of Biofoundries. *Nat. Commun.* **2019**, *10*, 2040. [[CrossRef](#)] [[PubMed](#)]
- [229] Matzko, R.; Konur, S. Technologies for Design-Build-Test-Learn Automation and Computational Modelling across the Synthetic Biology Workflow: A Review. *Netw. Model. Anal. Health Inform. Bioinform.* **2024**, *13*, 22. [[CrossRef](#)]
- [230] Goshisht, M.K. Machine Learning and Deep Learning in Synthetic Biology: Key Architectures, Applications, and Challenges. *ACS Omega* **2024**, *9*, 9921–9945. [[CrossRef](#)]
- [231] Holowko, M.B.; Frow, E.K.; Reid, J.C.; Rourke, M.; Vickers, C.E. Building a Biofoundry. *Synth. Biol.* **2021**, *6*, ysaa026. [[CrossRef](#)] [[PubMed](#)]
- [232] Bultelle, M.; Casas, A.; Kitney, R. Engineering Biology and Automation—Replicability as a Design Principle. *Eng. Biol.* **2024**, *8*, 53–68. [[CrossRef](#)] [[PubMed](#)]
- [233] Emmert-Streib, F.; Yli-Harja, O. What Is a Digital Twin? Experimental Design for a Data-Centric Machine Learning Perspective in Health. *Int. J. Mol. Sci.* **2022**, *23*, 13149. [[CrossRef](#)]
- [234] McLaughlin, J.A.; Beal, J.; Mısırlı, G.; Grünberg, R.; Bartley, B.A.; Scott-Brown, J.; Vaidyanathan, P.; Fontanarrosa, P.; Oberortner, E.; Wipat, A.; et al. The Synthetic Biology Open Language (SBOL) Version 3: Simplified Data Exchange for Bioengineering. *Front. Bioeng. Biotechnol.* **2020**, *8*, 1009. [[CrossRef](#)]
- [235] Gierend, K.; Krüger, F.; Genehr, S.; Hartmann, F.; Siegel, F.; Waltemath, D.; Ganslandt, T.; Zeleke, A.A. Provenance Information for Biomedical Data and Workflows: Scoping Review. *J. Med. Internet Res.* **2024**, *26*, e51297. [[CrossRef](#)]
- [236] Singh, N.; Lane, S.; Yu, T.; Lu, J.; Ramos, A.; Cui, H.; Zhao, H. A Generalized Platform for Artificial Intelligence-Powered Autonomous Enzyme Engineering. *Nat. Commun.* **2025**, *16*, 5648. [[CrossRef](#)]
- [237] Rapp, J.T.; Bremer, B.J.; Romero, P.A. Self-Driving Laboratories to Autonomously Navigate the Protein Fitness Landscape. *Nat. Chem. Eng.* **2024**, *1*, 97–107. [[CrossRef](#)]
- [238] Martin, H.G.; Radivojevic, T.; Zucker, J.; Bouchard, K.; Sustarich, J.; Peisert, S.; Arnold, D.; Hillson, N.; Babnigg, G.; Marti, J.M.; et al. Perspectives for Self-Driving Labs in Synthetic Biology. *Curr. Opin. Biotechnol.* **2023**, *79*, 102881. [[CrossRef](#)] [[PubMed](#)]

- [239] Melocchi, A.; Schmittlein, B.; Sadhu, S.; Nayak, S.; Lares, A.; Uboldi, M.; Zema, L.; di Robilant, B.N.; Feldman, S.A.; Esensten, J.H. Automated Manufacturing of Cell Therapies. *J. Control. Release* **2025**, *381*, 113561. [[CrossRef](#)]
- [240] Wang, B.; Chen, R.Q.; Li, J.; Roy, K. Interfacing Data Science with Cell Therapy Manufacturing: Where We Are and Where We Need to Be. *Cyotherapy* **2024**, *26*, 967–979. [[CrossRef](#)] [[PubMed](#)]
- [241] Gemmati, D.; Longo, G.; Gallo, I.; Silva, J.A.; Secchiero, P.; Zauli, G.; Hanau, S.; Passaro, A.; Pellegatti, P.; Pizzicotti, S.; et al. 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.* **2022**, *13*, 1028081. [[CrossRef](#)]
- [242] Morgan, R.A.; Chinnasamy, N.; Abate-Daga, D.; Gros, A.; Robbins, P.F.; Zheng, Z.; Dudley, M.E.; Feldman, S.A.; Yang, J.C.; Sherry, R.M.; et al. Cancer Regression and Neurological Toxicity Following Anti-MAGE-A3 TCR Gene Therapy. *J. Immunother.* **2013**, *36*, 133–151. [[CrossRef](#)]
- [243] Birnbaum, M.E.; Mendoza, J.L.; Sethi, D.K.; Dong, S.; Glanville, J.; Dobbins, J.; Özkan, E.; Davis, M.M.; Wucherpfennig, K.W.; Garcia, K.C. Deconstructing the Peptide-MHC Specificity of T Cell Recognition. *Cell* **2014**, *157*, 1073–1087. [[CrossRef](#)]
- [244] Cai, Y.; Chen, R.; Gao, S.; Li, W.; Liu, Y.; Su, G.; Song, M.; Jiang, M.; Jiang, C.; Zhang, X. Artificial Intelligence Applied in Neoantigen Identification Facilitates Personalized Cancer Immunotherapy. *Front. Oncol.* **2022**, *12*, 1054231. [[CrossRef](#)] [[PubMed](#)]
- [245] Zou, H.; Liu, W.; Wang, X.; Wang, Y.; Wang, C.; Qiu, C.; Liu, H.; Shan, D.; Xie, T.; Huang, W.; et al. 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* **2024**, *12*, e008450. [[CrossRef](#)] [[PubMed](#)]
- [246] Bottini, M.; Ryu, S.-J.; Terander, A.E.; Voglis, S.; Maldaner, N.; Bellut, D.; Regli, L.; Serra, C.; Staartjes, V.E. The Ever-Evolving Regulatory Landscape Concerning Development and Clinical Application of Machine Intelligence: Practical Consequences for Spine Artificial Intelligence Research. *Neurospine* **2025**, *22*, 134–143. [[CrossRef](#)] [[PubMed](#)]
- [247] Hersh, W.; Hollis, K.F. Results and Implications for Generative AI in a Large Introductory Biomedical and Health Informatics Course. *npj Digit. Med.* **2024**, *7*, 247. [[CrossRef](#)]
- [248] Chao, R.; Mishra, S.; Si, T.; Zhao, H. Engineering Biological Systems Using Automated Biofoundries. *Metab. Eng.* **2017**, *42*, 98–108. [[CrossRef](#)]
- [249] Nettleton, D.F.; Marí-Buyé, N.; Marti-Soler, H.; Egan, J.R.; Hort, S.; Horna, D.; Costa, M.; Benítez-Cano, E.V.; Goldrick, S.; Rafiq, Q.A.; et al. Smart Sensor Control and Monitoring of an Automated Cell Expansion Process. *Sensors* **2023**, *23*, 9676. [[CrossRef](#)]
- [250] Cheng, F.; Xie, W.; Zheng, H. Digital Twin Calibration for Biological System-of-Systems: Cell Culture Manufacturing Process. 2024 Winter Simulation Conference (WSC), United States, 2024; pp. 323–334. [[CrossRef](#)]
- [251] Hie, B.L.; Shanker, V.R.; Xu, D.; Bruun, T.U.J.; Weidenbacher, P.A.; Tang, S.; Wu, W.; Pak, J.E.; Kim, P.S. Efficient Evolution of Human Antibodies from General Protein Language Models. *Nat. Biotechnol.* **2024**, *42*, 275–283. [[CrossRef](#)]
- [252] Jonny, Sitepu, E.C.; Nidom, C.A.; Wirjopranoto, S.; Sudiana, I.K.; Ansori, A.N.M.; Putranto, T.A. Ex Vivo-Generated Tolerogenic Dendritic Cells: Hope for a Definitive Therapy of Autoimmune Diseases. *Curr. Issues Mol. Biol.* **2024**, *46*, 4035–4048. [[CrossRef](#)]
- [253] Stucchi, A.; Maspes, F.; Montee-Rodrigues, E.; Foustieri, G. Engineered Treg Cells: The Heir to the Throne of Immunotherapy. *J. Autoimmun.* **2024**, *144*, 102986. [[CrossRef](#)]
- [254] Wang, K.; Song, B.; Zhu, Y.; Dang, J.; Wang, T.; Song, Y.; Shi, Y.; You, S.; Li, S.; Yu, Z.; et al. Peripheral Nerve-Derived CSF1 Induces BMP2 Expression in Macrophages to Promote Nerve Regeneration and Wound Healing. *npj Regen. Med.* **2024**, *9*, 35. [[CrossRef](#)] [[PubMed](#)]
- [255] O'Brien, H.; Salm, M.; Morton, L.T.; Szukszto, M.; O'Farrell, F.; Boulton, C.; King, L.; Bola, S.K.; Becker, P.D.; Craig, A.; et al. A Modular Protein Language Modelling Approach to Immunogenicity Prediction. *PLoS Comput. Biol.* **2024**, *20*, e1012511. [[CrossRef](#)]
- [256] Alowidi, N.; Ali, R.; Sadaqah, M.; Naemi, F.M.A. Advancing Kidney Transplantation: A Machine Learning Approach to Enhance Donor–Recipient Matching. *Diagnostics* **2024**, *14*, 2119. [[CrossRef](#)]
- [257] Weimer, E.T.; Newhall, K.A. Machine Learning Enhanced Immunologic Risk Assessments for Solid Organ Transplantation. *Sci. Rep.* **2025**, *15*, 7943. [[CrossRef](#)] [[PubMed](#)]
- [258] Alsaafeen, B.H.; Ali, B.R.; Elkord, E. Combinational Therapeutic Strategies to Overcome Resistance to Immune Checkpoint Inhibitors. *Front. Immunol.* **2025**, *16*, 1546717. [[CrossRef](#)]
- [259] Liu, Z.; Zhang, J.; Hong, L.; Nie, Q.; Sun, X. Multiscale Mathematical Model-Informed Reinforcement Learning Optimizes Combination Treatment Scheduling in Glioblastoma Evolution. *Sci. Adv.* **2025**, *11*, eadv3316. [[CrossRef](#)] [[PubMed](#)]
- [260] Prelaj, A.; Miskovic, V.; Zanitti, M.; Trovo, F.; Genova, C.; Viscardi, G.; Rebuzzi, S.; Mazzeo, L.; Provenzano, L.; Kosta, S.; et al. Artificial Intelligence for Predictive Biomarker Discovery in Immuno-Oncology: A Systematic Review. *Ann. Oncol.* **2024**, *35*, 29–65. [[CrossRef](#)]
- [261] Ahlquist, K.D.; Sugden, L.A.; Ramachandran, S. Enabling Interpretable Machine Learning for Biologi-

- cal Data with Reliability Scores. *PLoS Comput. Biol.* **2023**, *19*, e1011175. [[CrossRef](#)]
- [262] Shiferaw, K.B.; Roloff, M.; Balaur, I.; Welter, D.; Waltemath, D.; Zeleke, A.A. Guidelines and Standard Frameworks for Artificial Intelligence in Medicine: A Systematic Review. *JAMIA Open* **2025**, *8*, ooae155. [[CrossRef](#)]
- [263] Vithlani, J.; Hawksworth, C.; Elvidge, J.; Ayiku, L.; Dawoud, D. Economic Evaluations of Artificial Intelligence-Based Healthcare Interventions: A Systematic Literature Review of Best Practices in Their Conduct and Reporting. *Front. Pharmacol.* **2023**, *14*, 1220950. [[CrossRef](#)]
- [264] Pannu, J.; Bloomfield, D.; MacKnight, R.; Hanke, M.S.; Zhu, A.; Gomes, G.; Cicero, A.; Inglesby, T.V. Dual-Use Capabilities of Concern of Biological AI Models. *PLoS Comput. Biol.* **2025**, *21*, e1012975. [[CrossRef](#)]
- [265] Bloomfield, D.; Pannu, J.; Zhu, A.W.; Ng, M.Y.; Lewis, A.; Bendavid, E.; Asch, S.M.; Hernandez-Boussard, T.; Cicero, A.; Inglesby, T. AI and Biosecurity: The Need for Governance. *Science* **2024**, *385*, 831–833. [[CrossRef](#)] [[PubMed](#)]
- [266] Derraz, B.; Breda, G.; Kaempf, C.; Baenke, F.; Cotte, F.; Reiche, K.; Köhl, U.; Kather, J.N.; Eskenazy, D.; Gilbert, S. New Regulatory Thinking Is Needed for AI-Based Personalised Drug and Cell Therapies in Precision Oncology. *npj Precis. Oncol.* **2024**, *8*, 23. [[CrossRef](#)]
- [267] Groff-Vindman, C.S.; Trump, B.D.; Cummings, C.L.; Smith, M.; Titus, A.J.; Oye, K.; Prado, V.; Turmus, E.; Linkov, I. The Convergence of AI and Synthetic Biology: The Looming Deluge. *npj Biomed. Innov.* **2025**, *2*, 20. [[CrossRef](#)]
- [268] Undheim, T.A. The Whack-a-Mole Governance Challenge for AI-Enabled Synthetic Biology: Literature Review and Emerging Frameworks. *Front. Bioeng. Biotechnol.* **2024**, *12*, 1359768. [[CrossRef](#)] [[PubMed](#)]
- [269] Wu, J.; Koelzer, V.H. Towards Generative Digital Twins in Biomedical Research. *Comput. Struct. Biotechnol. J.* **2024**, *23*, 3481–3488. [[CrossRef](#)] [[PubMed](#)]
- [270] Liu, S.; Guo, L.R. Data Ownership in the AI-Powered Integrative Health Care Landscape. *JMIR Med. Inform.* **2024**, *12*, e57754. [[CrossRef](#)]
- [271] Welzel, C.; Ostermann, M.; Smith, H.L.; Minssen, T.; Kirsten, T.; Gilbert, S. Enabling Secure and Self-Determined Health Data Sharing and Consent Management. *npj Digit. Med.* **2025**, *8*, 560. [[CrossRef](#)]
- [272] Balsa-Canto, E.; Campo-Manzanares, N.; Moimenta, A.R.; Roudaut, G.; Troitiño-Jordedo, D. Quantifying and Managing Uncertainty in Systems Biology: Mechanistic and Data-Driven Models. *Curr. Opin. Syst. Biol.* **2025**, *42*, 100557. [[CrossRef](#)]
- [273] Zhou, Z.; Hu, M.; Salcedo, M.; Gravel, N.; Yeung, W.; Venkat, A.; Guo, D.; Zhang, J.; Kannan, N.; Li, S. XAI Meets Biology: A Comprehensive Review of Explainable AI in Bioinformatics Applications. *arXiv* **2023**, arXiv:2312.06082. [[CrossRef](#)]
- [274] Singh, R.; Paxton, M.; Auclair, J. Regulating the AI-Enabled Ecosystem for Human Therapeutics. *Commun. Med.* **2025**, *5*, 181. [[CrossRef](#)]
- [275] Niazi, S.K. Regulatory Perspectives for AI/ML Implementation in Pharmaceutical GMP Environments. *Pharmaceuticals* **2025**, *18*, 901. [[CrossRef](#)]
- [276] de Lima, R.C.; Quaresma, J.A.S. Emerging Technologies Transforming the Future of Global Biosecurity. *Front. Digit. Health* **2025**, *7*, 1622123. [[CrossRef](#)]
- [277] Kumar, D.; Malin, B.A.; Vishwanatha, J.K.; Wu, L.; Hedges, J.R. AI in Biomedicine—A Forward-Looking Perspective on Health Equity. *Int. J. Environ. Res. Public Health* **2024**, *21*, 1642. [[CrossRef](#)]
- [278] Johnson, K.B.; Horn, I.B.; Horvitz, E. Pursuing Equity with Artificial Intelligence in Health Care. *JAMA Health Forum* **2025**, *6*, e245031. [[CrossRef](#)]
- [279] Krenn, M.; Pollice, R.; Guo, S.Y.; Aldeghi, M.; Cervera-Lierta, A.; Friederich, P.; Gomes, G.d.P.; Häse, F.; Jinich, A.; Nigam, A.; et al. On Scientific Understanding with Artificial Intelligence. *Nat. Rev. Phys.* **2022**, *4*, 761–769. [[CrossRef](#)] [[PubMed](#)]
- [280] Alvarado, R. AI as an Epistemic Technology. *Sci. Eng. Ethics* **2023**, *29*, 32. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher's Note: The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the publisher or editors. The publisher and editors assume no responsibility for any injury or damage resulting from the use of information contained herein.