

# BIOCHEMICAL HALLMARKS OF SKIN DISORDERS AND THE EMERGING ROLE OF ARTIFICIAL INTELLIGENCE IN THERAPEUTICS: A COMPREHENSIVE REVIEW

Samra Wahid<sup>1</sup>, Hamza Rafeeq<sup>\*2</sup>

<sup>1</sup>Department of Dermal Science, Riphah International University Faisalabad Campus, Faisalabad, Punjab Pakistan. 44000.

<sup>\*2</sup>Department of Biochemistry, Riphah International University Faisalabad Campus, Faisalabad, Punjab Pakistan. 44000

<sup>\*2</sup>[hamza.rafeeq@riphahfsd.edu.pk](mailto:hamza.rafeeq@riphahfsd.edu.pk)

DOI: <https://doi.org/10.5281/zenodo.20622795>

## Keywords

skin disorders; biochemistry; artificial intelligence; machine learning; deep learning; precision dermatology; biomarkers; multi-omics; drug discovery; digital therapeutics

## Article History

Received: 03 April 2026

Accepted: 15 May 2026

Published: 30 May 2026

Copyright @Author

Corresponding Author: \*

Hamza Rafeeq

## Abstract

Skin disorders affect close to one-third of the global population and impose a disability burden that rivals many internal diseases, yet they remain conspicuously under-prioritised in research and health policy. Their biochemical substrate is remarkably heterogeneous, spanning chronic cytokine-driven inflammation in psoriasis and atopic dermatitis, dysregulated sebaceous lipid metabolism and microbial dysbiosis in acne, disordered melanogenesis in vitiligo and melasma, and oncogenic signalling with metabolic reprogramming in keratinocyte and melanocytic cancers. Conventional diagnostics and therapeutics are frequently limited by late detection, the subjective nature of visual and histological assessment, marked inter-individual variability in drug response, and an incomplete understanding of pathway crosstalk. Artificial intelligence (AI), and in particular machine learning and deep learning, offers a route to integrate multimodal biochemical, imaging, and clinical data at a scale and granularity beyond human capacity. This review first synthesises the biochemical hallmarks of the major skin diseases and then critically examines how AI is being deployed to (i) discover and prioritise molecular biomarkers from genomic, transcriptomic, proteomic, and metabolomic datasets; (ii) raise diagnostic accuracy above that of conventional histology and unaided clinicians; (iii) predict individual responses to biologics and small-molecule inhibitors and personalise treatment sequences and (iv) accelerate drug discovery and repurposing for both common and rare dermatoses. We close by appraising the principal obstacles to clinical translation - data heterogeneity and the absence of standardised biochemical ontologies, the opacity of high-capacity models, algorithmic bias across skin phototypes, and regulatory and privacy concerns and outline a path towards an explainable, multimodal, and equitable digital dermatology that genuinely complements biochemical knowledge rather than displacing it.

## 1. Introduction

The skin is the largest organ of the human body and its most exposed interface with the external environment, simultaneously functioning as a

physical barrier, an immunological sentinel, a metabolic factory, and an instrument of thermoregulation and sensation. Disorders of the skin are correspondingly diverse, and they are

extraordinarily common. Analyses derived from the Global Burden of Disease programme rank skin and subcutaneous diseases among the leading causes of non-fatal disease burden worldwide, with collective prevalence estimates approaching one-third of the population at any given time [1,3]. The morbidity they generate is far from trivial: chronic inflammatory dermatoses are associated with disfigurement, pruritus, sleep disturbance, depression, social stigmatisation, and substantial loss of productivity, while cutaneous malignancies, melanoma in particular, contribute meaningfully to cancer mortality [2,4,22]. Despite this, dermatological conditions have historically attracted research investment that is disproportionately small relative to their burden, and access to specialist care remains deeply inequitable across low- and middle-income settings [2,4].

Underlying this clinical heterogeneity is an equally heterogeneous biochemistry. Psoriasis and atopic dermatitis are sustained by self-amplifying cytokine circuits, oxidative stress, and disrupted lipid and barrier biology [6,7,9,11]. Acne arises from the convergence of androgen-driven sebaceous hyperactivity, altered sebum lipid composition, follicular hyperkeratinisation, and the metabolic activity of *Cutibacterium acnes* [13,15]. Pigmentary disorders reflect perturbed melanocyte biology and tyrosinase-dependent melanin synthesis, whether through autoimmune melanocyte destruction in vitiligo or hormone- and light-driven overproduction in melasma [17,19,20]. Cutaneous cancers are driven by characteristic somatic mutations in the mitogen-activated protein kinase (MAPK), PTEN-PI3K, and p53 pathways, accompanied by a metabolic shift that favours proliferation and immune evasion [21,23,24]. A recurring theme across all of these conditions is that no single molecule acts in isolation; rather, dense, redundant, and often nonlinear networks of signalling determine phenotype and therapeutic response.

Conventional approaches to diagnosis and treatment struggle precisely with this complexity. Visual inspection and dermoscopy are operator-dependent and vary widely in accuracy; histopathology, although the reference standard

for many conditions, is qualitative and subject to inter-observer disagreement; and biochemical assays typically interrogate one analyte at a time, providing a narrow keyhole onto a panoramic network. Therapeutically, even the most effective modern biologics produce adequate responses in only a fraction of patients, and the choice among an expanding menu of targeted agents remains largely empirical [64,65]. The result is a costly cycle of trial and error, delayed control of disease, and avoidable adverse effects.

Artificial intelligence has emerged as a candidate solution to several of these problems. Deep neural networks can learn hierarchical representations directly from raw data, dispensing with hand-crafted features and capturing the nonlinear interactions that defeat classical statistics [28]. The watershed demonstration in dermatology was the classification of skin cancer at dermatologist level by a convolutional neural network trained on a large image corpus [29], and the field has since expanded to encompass molecular profiling, electronic health record mining, and therapeutic optimisation [30,31,38]. Critically, AI is not merely an image classifier: its greatest promise lies in fusing imaging with the biochemical and multi-omic data that define disease mechanism, producing predictions that are both accurate and mechanistically interpretable.

The scope of this review is deliberately integrative rather than exhaustive. We focus on the biochemical mechanisms that are both well characterised and clinically actionable, and on the AI methods that have either reached clinical evaluation or demonstrated clear translational potential. We draw selectively on the literature of immunology, lipidomics, melanocyte biology, and cancer genetics on the biochemical side, and on computer vision, multi-omic integration, and clinical informatics on the computational side, with the aim of identifying where the two disciplines most productively intersect. Where the evidence is strong we say so; where claims of performance rest on narrow benchmarks or retrospective data we flag the limitation explicitly, because the credibility of the field depends on distinguishing genuine clinical value from impressive but fragile accuracy figures.

Throughout, biochemical mechanism is treated as the connective tissue that makes AI predictions interpretable and trustworthy, and the recurring argument is that the marriage of molecular insight with computational power, rather than either alone, defines the frontier of modern dermatology.

This review is structured to bridge these two worlds. We first set out, disorder by disorder, the biochemical hallmarks that define the major dermatoses and the pathways currently exploited by targeted therapy (Section 2). We then describe the principal families of AI techniques and how they map onto biochemical and imaging data (Section 3), before examining their application to diagnosis and biochemical profiling (Section 4), to drug discovery and personalised treatment (Section 5), and to monitoring and the pursuit of cure in real-world settings (Section 6). We conclude with a candid assessment of the scientific, technical, ethical, and regulatory barriers to adoption, and a forward-looking perspective on an integrated digital dermatology (Section 7). Throughout, our aim is to treat AI not as a replacement for biochemical understanding but as an amplifier of it.

## 2. Biochemical Basis of Major Skin Disorders

A coherent application of AI to dermatology must rest on a clear account of what, biochemically, is being measured and modelled. This section surveys the molecular pathology of four disease groups that together account for the majority of the dermatological disease burden and that recur throughout the AI literature. Table 1 consolidates the key pathways, diagnostic biomarkers, molecular targets, and exemplar drugs for these disorders.

### 2.1 Inflammatory disorders: psoriasis and atopic dermatitis

Psoriasis is the prototypical T-cell-mediated inflammatory dermatosis. Its pathogenesis is now understood as a feed-forward loop in which activated dendritic cells secrete interleukin (IL)-23 and IL-12, driving the differentiation and maintenance of T helper 17 (Th17) and Th1 cells [5,7]. The Th17 effector cytokines IL-17A, IL-17F,

and IL-22, together with tumour necrosis factor (TNF)-alpha, act on keratinocytes to induce hyperproliferation, defective terminal differentiation, and the production of antimicrobial peptides and chemokines that recruit further inflammatory cells, thereby sustaining the circuit [7,8]. The clinical efficacy of agents targeting individual nodes of this network - anti-TNF agents, the anti-IL-12/23 p40 antibody, anti-IL-23 p19 antibodies, and the anti-IL-17 antibodies - provides striking confirmation of its centrality [8,61,62]. At the tissue level, psoriatic lesions display elevated reactive oxygen species, perturbed polyamine and arachidonic acid metabolism, and an altered lipidomic signature, all of which constitute candidate biomarkers [26].

Atopic dermatitis, by contrast, is driven predominantly by a type 2 immune response, although it is increasingly recognised as a heterogeneous, endotype-rich disease. Two intertwined defects define it: impairment of the epidermal barrier and dysregulated immunity [9,11]. Loss-of-function mutations in the gene encoding filaggrin, a structural protein essential for corneocyte formation and the generation of natural moisturising factor, are the strongest known genetic risk factor and exemplify how a single biochemical lesion compromises barrier integrity, raises skin pH, and facilitates allergen penetration and microbial colonisation [10]. The ensuing immune response is dominated by IL-4 and IL-13, which further suppress filaggrin and antimicrobial peptide expression, alongside IL-31, a principal mediator of pruritus, and IL-22 and IL-33 [12]. The therapeutic success of dupilumab, which blocks the shared IL-4 receptor alpha subunit and thereby neutralises both IL-4 and IL-13 signalling, validates this axis [63]. Stratum corneum lipid abnormalities - reduced long-chain ceramides and altered free fatty acid profiles - are consistent, quantifiable biochemical hallmarks and attractive targets for both diagnosis and barrier-directed therapy [27].

From a data-science standpoint, these two conditions are instructive because they can be distinguished not only by morphology but by their cytokine and transcriptomic signatures: the Th17/IL-17 axis dominates psoriasis whereas the

Th2/IL-13 axis dominates atopic dermatitis, with corresponding differences in barrier gene expression [7,12,52]. This molecular separability is precisely what makes them tractable targets for supervised machine learning on biochemical inputs.

Beyond their cytokine-defining axes, both diseases share biochemical themes that are increasingly recognised as drivers of chronicity and as candidate biomarkers. Oxidative stress, reflecting an imbalance between reactive oxygen species generation and antioxidant defence, is a feature of psoriatic and eczematous skin alike, contributing to lipid peroxidation, protein and DNA damage, and the amplification of inflammatory signalling [26]. Lipid mediators derived from arachidonic acid - prostaglandins, leukotrienes, and related eicosanoids - modulate vascular and inflammatory responses, while disturbances in epidermal sphingolipid and ceramide metabolism compromise the permeability barrier and feed back onto immune activation [27]. These metabolic and lipidomic alterations are quantifiable by mass spectrometry and constitute a rich, high-dimensional feature space that is well suited to machine-learning analysis but poorly captured by conventional single-analyte assays. The dynamic, fluctuating nature of these biochemical signals over the course of disease flares and remissions further argues for the continuous, computationally interpreted monitoring discussed in Section 6.

## 2.2 Acne and sebaceous gland disorders

Acne vulgaris arises within the pilosebaceous unit through the interaction of four classical pathogenic factors: increased and qualitatively altered sebum production, abnormal follicular keratinisation, colonisation and metabolic activity of *Cutibacterium acnes*, and inflammation [13,14]. Androgens stimulate sebocyte proliferation and lipogenesis, and the resulting sebum is enriched in triglycerides, wax esters, and squalene; oxidation of squalene and depletion of antioxidants such as vitamin E render the lipid environment pro-inflammatory and comedogenic [15]. Specific phylotypes of *C. acnes*, rather than simple bacterial load, appear to drive disease by

activating innate immune sensors and metabolising sebum lipids into pro-inflammatory free fatty acids [15]. Insulin and insulin-like growth factor 1 signalling, modulated by diet, amplify sebaceous lipogenesis and androgen activity, providing a biochemical rationale for the association between high-glycaemic-load diets and acne severity [16]. The mainstays of severe acne therapy - systemic isotretinoin, which profoundly suppresses sebaceous gland activity, and hormonal agents - act directly on these biochemical drivers, while the sebaceous gland itself is increasingly viewed as an endocrine and immunocompetent organ whose lipidome offers a rich substrate for quantitative profiling [14,15].

From an analytical perspective, acne is notable for the richness and accessibility of its biochemical readouts. The skin surface lipidome can be sampled non-invasively and quantified by mass spectrometry, yielding compositional signatures that correlate with disease activity and with response to sebosuppressive therapy [15]. The cutaneous microbiome, and specifically the relative abundance of pathogenic versus commensal *C. acnes* phylotypes, can be characterised by high-throughput sequencing, generating high-dimensional feature sets well suited to machine-learning classification [13]. Together these lipidomic and microbiomic data offer objective, continuous endpoints to supplement the coarse and observer-dependent lesion counts on which acne severity grading has traditionally relied. This abundance of quantifiable molecular features positions acne as a particularly tractable target for the integrative, biochemistry-informed AI approaches developed in later sections, where surface lipid composition and microbial profiles can be modelled jointly with clinical imaging to stratify severity and predict therapeutic response [16].

## 2.3 Pigmentation disorders: vitiligo and melasma

Pigmentary disorders reflect disturbances of melanocyte number, function, or melanin distribution. Melanin synthesis proceeds within melanosomes through a tyrosinase-catalysed cascade that converts tyrosine to dopaquinone and

ultimately to eumelanin and pheomelanin; tyrosinase, tyrosinase-related protein 1, and dopachrome tautomerase are the rate-limiting enzymes and the principal targets of depigmenting therapy [19]. Vitiligo is an autoimmune disease in which melanocytes are selectively destroyed. Oxidative stress within melanocytes generates damage-associated molecular patterns and stress proteins that, in genetically susceptible individuals, trigger an interferon-gamma-driven, CXCL9/CXCL10-CXCR3 chemokine axis recruiting cytotoxic CD8<sup>+</sup> T cells [17,18]. The recent clinical success of Janus kinase (JAK) inhibitors in repigmenting vitiligo directly validates this interferon-JAK-signal transducer and activator of transcription pathway as a therapeutic target and illustrates how biochemical insight translates into treatment [18]. Melasma, in contrast, is a disorder of melanocyte hyperfunction in which ultraviolet radiation, visible light, hormones, and vascular and senescence-associated factors converge to upregulate melanogenesis; it is increasingly conceptualised as a photoaging phenomenon involving the dermal microenvironment as well as the epidermis [20]. These enzymatic and chemokine-based signatures are quantifiable and provide objective endpoints against which AI models of disease activity and treatment response can be trained.

#### **2.4 Skin cancers: melanoma and keratinocyte carcinomas**

Cutaneous malignancies are defined by characteristic genomic alterations that activate proliferative signalling and disable tumour-suppressor and DNA-repair mechanisms, frequently against a backdrop of ultraviolet-induced mutational signatures. In melanoma, activating mutations in BRAF (most commonly V600E) and NRAS constitutively switch on the MAPK pathway, while losses of PTEN and CDKN2A and mutations in TP53 and TERT promoter regions cooperate to drive progression [21,23,24]. The translation of the BRAF V600E discovery into BRAF and MEK inhibitor therapy transformed the management of advanced melanoma and remains a paradigm of biomarker-

driven precision oncology [25]. Basal cell carcinoma is overwhelmingly driven by aberrant Hedgehog pathway signalling, typically through PTCH1 loss or SMO activation, providing the rationale for Hedgehog pathway inhibitors, whereas cutaneous squamous cell carcinoma carries a very high ultraviolet-associated mutational burden with frequent TP53, NOTCH, and CDKN2A alterations [23]. Across these tumours, metabolic reprogramming - increased glycolytic flux, altered lipid and glutamine metabolism, and adaptation of the tumour microenvironment - supports rapid growth and immune evasion and offers both diagnostic biomarkers and emerging therapeutic vulnerabilities [23,24]. The wealth of genomic, transcriptomic, and histological data generated in oncology has made skin cancer the single most intensively studied application of AI in dermatology, as discussed in Section 4.

An important and rapidly expanding dimension of cutaneous oncology biochemistry concerns the tumour microenvironment and systemic biomarkers. The immune contexture of a melanoma - the density and functional state of tumour-infiltrating lymphocytes, the expression of immune-checkpoint ligands, and the metabolic conditioning of the microenvironment - shapes both prognosis and response to immunotherapy, and is increasingly profiled by multiplexed and spatial molecular techniques amenable to AI analysis [23,24]. In parallel, minimally invasive biomarkers measurable in blood, including circulating tumour DNA carrying characteristic BRAF or NRAS mutations and lactate dehydrogenase as a metabolic surrogate, offer the prospect of monitoring disease burden and treatment response longitudinally [24,25]. These quantitative, dynamic biochemical readouts are natural inputs for predictive models, and their integration with imaging and histology exemplifies the multimodal strategy that recurs throughout this review.

#### **2.5 Shared biochemical themes and the rationale for integrative analysis**

Although the disorders surveyed above are clinically and mechanistically distinct, several

biochemical themes recur across them and motivate an integrative, computational approach. First, all are governed by networks rather than single molecules: cytokine circuits, enzymatic cascades, and oncogenic signalling pathways exhibit extensive crosstalk, feedback, and redundancy that defeat reductionist, one-target-at-a-time analysis. Second, oxidative stress and altered lipid metabolism appear, in different guises, across inflammatory, pigmentary, and neoplastic skin disease, suggesting shared vulnerabilities and the possibility of cross-cutting biomarkers [26,27]. Third, the relevant biochemistry is dynamic, varying with disease

activity, treatment, and environmental exposure, so that static single-timepoint measurements capture only a fraction of the information available. Fourth, each disorder is increasingly resolvable into molecular endotypes that cut across traditional morphological categories and that carry therapeutic implications [50,52]. These features - network complexity, shared metabolic axes, temporal dynamics, and molecular heterogeneity - are precisely those that classical statistics handle poorly and that machine learning, given adequate data, handles well. They constitute the biochemical case for the AI methods examined in the remainder of this review.

**Table 1. Key biochemical pathways, diagnostic biomarkers, molecular targets, and exemplar therapies for the major skin disorders.**

Disorder	Key biochemical pathway(s)	Diagnostic biomarkers	Current molecular target(s)	Example drug(s)
Psoriasis	IL-23/Th17 axis; TNF-alpha; keratinocyte hyperproliferation; oxidative stress	IL-17A, IL-23, TNF-alpha; antimicrobial peptides; CRP	TNF-alpha; IL-12/23 p40; IL-23 p19; IL-17A/RA	Adalimumab; ustekinumab; guselkumab; secukinumab, ixekizumab
Atopic dermatitis	Type 2 immunity (IL-4/IL-13); barrier dysfunction; IL-31 (itch)	Filaggrin mutations; serum IgE; TARC/CCL17; reduced ceramides	IL-4R-alpha (IL-4/IL-13); IL-13; IL-31R; JAK	Dupilumab; tralokinumab; nemolizumab; upadacitinib
Acne vulgaris	Androgen-driven sebaceous lipogenesis; C. acnes metabolism; IGF-1/insulin; inflammation	Sebum lipid profile; squalene oxidation; C. acnes phylotype	Retinoid signalling; androgen receptor; sebaceous lipogenesis	Isotretinoin; topical retinoids; hormonal agents
Vitiligo / melasma	Melanocyte oxidative stress; IFN-gamma-CXCL9/10-CXCR3 (vitiligo); tyrosinase-driven melanogenesis (melasma)	Tyrosinase activity; CXCL10; melanin indices	JAK-STAT/IFN axis (vitiligo); tyrosinase (melasma)	Ruxolitinib (topical); hydroquinone, tranexamic acid
Melanoma / keratinocyte cancers	MAPK (BRAF/NRAS); PI3K-PTEN; p53; Hedgehog (BCC); UV mutational	BRAF V600E; NRAS; TERT promoter; PTCH1/SMO (BCC)	BRAF; MEK; SMO; immune checkpoints	Vemurafenib + cobimetinib; vismodegib; checkpoint inhibitors

Disorder	Key biochemical pathway(s)	Diagnostic biomarkers	Current molecular target(s)	Example drug(s)
	signatures; metabolic reprogramming			

### 3. Artificial Intelligence Techniques Used in Dermatology and Biochemistry

AI in medicine is not a single technology but a family of computational methods that learn patterns from data. Their relevance to dermatology depends on the type of data being analysed - images, structured biochemical measurements, free text, or sequential treatment records - and on whether the task is classification, regression, clustering, or sequential decision-making. Figure 1 schematises how heterogeneous biochemical inputs (genomics, proteomics, metabolomics) are channelled through different model classes to yield clinical outputs. This section reviews the principal techniques and their characteristic strengths and limitations.

#### 3.1 Classical machine learning for biomarker discovery

Classical, or shallow, machine learning remains the workhorse for analysing structured biochemical data, particularly when sample sizes are modest relative to the number of measured features - the high-dimensional, low-sample regime typical of omics studies [41]. Random forests, ensembles of decision trees trained on bootstrapped samples and random feature subsets, are valued for their robustness to noise, their ability to model nonlinear interactions, and the intrinsic feature-importance rankings they provide, which lend themselves naturally to biomarker prioritisation [39]. Support vector machines, which seek the maximally separating hyperplane in a transformed feature space, perform well on high-dimensional data such as gene expression matrices and were among the earliest methods applied to molecular classification of disease [40,41]. Regularised regression methods, gradient-boosted trees, and penalised approaches that enforce sparsity are widely used to identify compact, interpretable panels of transcripts, proteins, or metabolites that

discriminate disease states or predict outcomes. Because these models operate on explicitly defined features, their outputs can be related directly to underlying biochemistry, an interpretability advantage that remains important in mechanistic and translational research [41].

The principal methodological challenge in this setting is the curse of dimensionality: omic datasets routinely measure thousands to tens of thousands of features in cohorts of only tens or hundreds of patients, creating ample opportunity for models to fit noise rather than signal. Robust biomarker discovery therefore depends as much on disciplined feature selection, dimensionality reduction, and validation as on the choice of algorithm. Techniques such as recursive feature elimination, penalised regression, and stability selection identify compact and reproducible feature panels, while unsupervised projection methods condense high-dimensional data into interpretable low-dimensional structure. Rigorous, nested cross-validation and, critically, validation on fully independent external cohorts are essential to guard against the optimistic performance estimates that plague underpowered studies; a biomarker that is not reproducible across cohorts and platforms is of little clinical value [41]. These statistical safeguards are not optional refinements but the difference between a genuine biochemical discovery and a spurious artefact of small samples, a theme to which Section 4 returns in the context of diagnostic models.

#### 3.2 Deep learning for image and histopathology analysis

Deep learning, and convolutional neural networks (CNNs) in particular, has driven the most visible advances in dermatological AI. CNNs apply hierarchies of learnable filters that progressively extract edges, textures, and increasingly abstract morphological features directly from pixel data, eliminating the need for hand-engineered

descriptors [28,42]. The broad utility of these methods across radiology, pathology, and dermatology has been documented in comprehensive surveys of deep learning in medical image analysis [32]. Architectural innovations such as deep residual networks, which use skip connections to enable the training of very deep models, underpin most modern dermatological classifiers [43], while encoder-decoder architectures such as U-Net are standard for the segmentation of lesions and histological structures [44]. In dermatology these networks have been applied to dermoscopic images, clinical photographs, and whole-slide histopathology, in each case learning representations that, given sufficient and representative training data, can match or exceed expert performance on narrowly defined tasks [29,33,37]. Transfer learning, in which a network pre-trained on a large generic image corpus is fine-tuned on a smaller dermatological dataset, is near-universal practice and partially mitigates the chronic scarcity of labelled medical images. The same convolutional machinery is increasingly applied to spatially resolved biochemical data, such as imaging mass spectrometry and multiplexed immunofluorescence, blurring the historical boundary between image analysis and molecular profiling.

A persistent constraint on deep learning in dermatology is the scarcity of large, well-annotated, and demographically representative datasets, since expert labelling is costly and histological or molecular ground truth is not always available. Several strategies mitigate this. Data augmentation - synthetically expanding training sets through rotation, colour transformation, and related perturbations - improves robustness and reduces overfitting, while generative models can produce realistic synthetic images to enrich rare classes, albeit with the risk of introducing artefacts. Self-supervised and contrastive learning, which learn useful representations from unlabelled data before fine-tuning on small labelled sets, are particularly promising for medical domains where unlabelled images are abundant but annotations are not. These approaches partially address the data bottleneck, but none substitutes for genuinely

diverse, prospectively collected data, and their use must be accompanied by careful evaluation to ensure that synthetic or augmented data do not encode or amplify the biases discussed in Section 7 [79].

### 3.3 Natural language processing for biochemical and clinical text

An enormous fraction of biochemical and clinical knowledge exists only as unstructured text - in electronic health records, pathology reports, and the biomedical literature. Natural language processing (NLP), transformed in recent years by transformer-based language models, allows this text to be mined at scale [45]. Domain-adapted models pre-trained on biomedical corpora capture the specialised vocabulary of biochemistry and medicine and substantially outperform general-purpose models on tasks such as named-entity recognition, relation extraction, and literature-based discovery [46]. In dermatology, NLP can extract phenotypes, drug exposures, and laboratory values from records to construct cohorts for predictive modelling, surface candidate gene-disease or drug-target relationships from the literature, and automate the structuring of pathology reports [47]. Coupled with the structured outputs of laboratory and omics platforms, NLP enables the assembly of the large, richly annotated datasets on which downstream predictive models depend.

### 3.4 Reinforcement learning for treatment optimization

Whereas the foregoing methods address perception and prediction, reinforcement learning (RL) addresses sequential decision-making: how to choose a series of actions over time to maximise a long-term outcome [48]. In medicine, RL frames treatment as a sequence of interventions in which each clinical state informs the next therapeutic choice, with the goal of learning a policy that optimises disease control while minimising toxicity. Proof-of-concept work in critical care has shown that RL can derive treatment strategies associated with improved outcomes from retrospective data [49]. Although still nascent in dermatology, the framework is naturally suited to

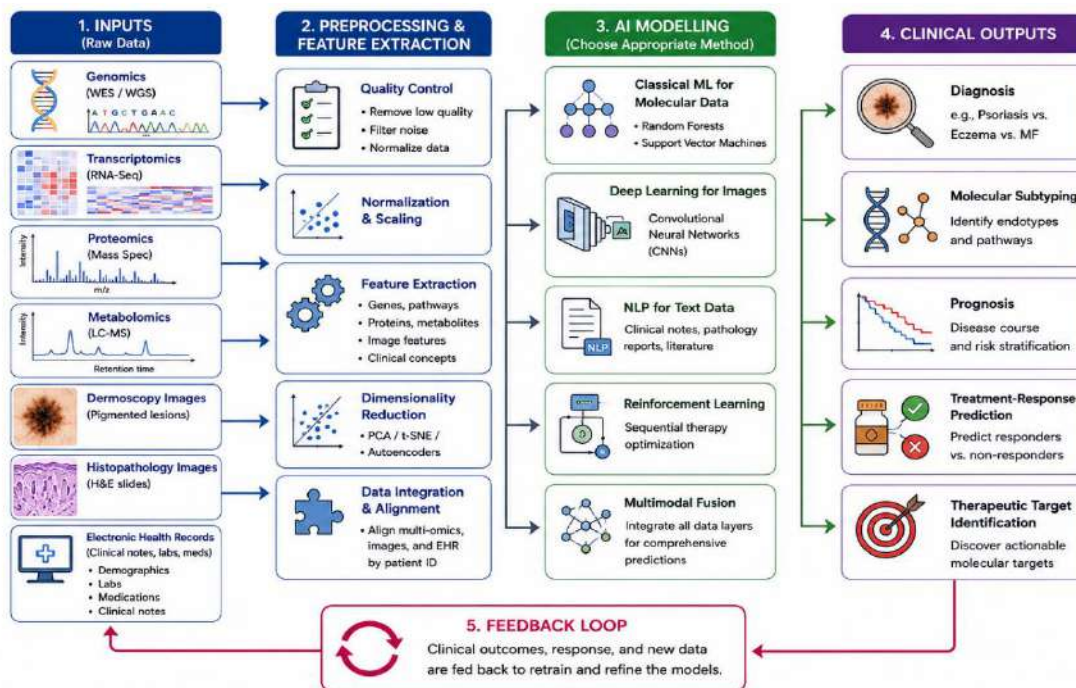
chronic diseases such as psoriasis and atopic dermatitis, where clinicians must iteratively titrate and switch among topical agents, phototherapy, and systemic biologics in response to evolving disease activity and biomarkers. The principal obstacles are the need for high-quality longitudinal data and the safety and ethical constraints on learning policies that will ultimately guide care, which together demand careful offline evaluation before any clinical deployment.

### 3.5 Multi-omics integration and the rise of multimodal models

Perhaps the most consequential trend is the move from single-data-type models towards integrative, multimodal architectures that combine genomics, transcriptomics, proteomics, metabolomics, imaging, and clinical variables within a unified framework [50]. Each omic layer captures a different stratum of biology, and their integration can reveal disease mechanisms and patient subgroups invisible to any single modality [50,90]. Methods range from early fusion, in which features from all modalities are concatenated before modelling, through intermediate fusion using learned joint representations, to late fusion of modality-specific predictions. The recent emergence of large foundation models and generalist multimodal architectures has accelerated this trajectory, raising the prospect of systems that reason jointly over a patient's images, molecular profiles, and records [89,90]. As

detailed in subsequent sections, it is this fusion of imaging with biochemistry, rather than image analysis alone, that distinguishes the most clinically promising dermatological AI.

Across all of these techniques, a set of methodological principles separates credible work from the merely impressive. Models must be evaluated on data held out from training, and ideally on external cohorts collected at different sites and times, because performance on internal test sets routinely overstates real-world accuracy [54,79]. The choice of metric matters: in the imbalanced datasets typical of dermatology, where malignant or rare cases are scarce, overall accuracy can be misleading and measures such as the area under the receiver-operating-characteristic curve, sensitivity at clinically relevant specificity, and calibration are more informative [53]. Equally, biochemical and imaging datasets are prone to confounding - scanner or assay batch effects, surgical skin markings, or site-specific labelling conventions - that models can exploit as shortcuts, achieving high benchmark scores while learning nothing biologically meaningful [83]. Rigorous reporting, pre-registered analysis plans, and independent validation are therefore not bureaucratic niceties but the conditions under which a model's claimed biochemical insight can be believed. These principles frame the applications surveyed in the remainder of this review.



**Figure 1.** Conceptual workflow of artificial intelligence in dermatological biochemistry. Raw biochemical and clinical inputs (genomics, transcriptomics, proteomics, metabolomics, dermoscopic and histopathological images, and electronic health records) undergo preprocessing and feature extraction, are modelled by an appropriate AI method (classical machine learning such as random forests and support vector machines for structured molecular data; convolutional neural networks for images; natural language processing for text; reinforcement learning for sequential therapy; and multimodal fusion across all layers), and yield clinical outputs spanning diagnosis, molecular subtyping, prognosis, treatment-response prediction, and therapeutic target identification. A feedback loop returns clinical outcomes to retrain and refine the models.

#### 4. AI for Diagnosis and Biochemical Profiling

Diagnosis is the domain in which dermatological AI has advanced furthest, and it is also the domain in which the integration of imaging with biochemistry is most clearly transformative. This section traces that progression from image-based classification, through molecular subtyping, to the combined analyses that increasingly define state-of-the-art systems. Table 2 compares representative AI diagnostic models by input type, biochemical features, and reported accuracy.

##### 4.1 From image classification to molecularly informed diagnosis

The foundational result was the demonstration that a single CNN, trained on a corpus of clinical and dermoscopic images, could classify keratinocyte carcinomas and melanomas at a level comparable to board-certified dermatologists [29].

Subsequent reader studies reinforced and extended this finding: a deep network was shown to outperform the majority of dermatologists in dermoscopic melanoma recognition, and large international comparisons confirmed that well-trained algorithms can equal or exceed average human accuracy on curated image sets [34,35,36]. Curated, publicly available datasets such as HAM10000 and the broader International Skin Imaging Collaboration archive have been instrumental, providing the labelled images on which most of these models are trained and benchmarked [33,55]. Models have since broadened beyond cancer to encompass differential diagnosis across scores of skin conditions from clinical photographs, approaching the accuracy of dermatologists for many common presentations [38,82].

Image-only diagnosis, however, has intrinsic limits. Two lesions that look alike may be biologically distinct, and morphologically ambiguous cases - precisely those for which a clinician most needs help - are where pixel data alone is least informative. The decisive advance is therefore the incorporation of biochemical and molecular information. Spatially resolved molecular techniques and the pairing of histological images with genomic or transcriptomic annotations allow models to learn associations between visual phenotype and underlying biochemistry, yielding diagnoses that are both more accurate and more mechanistically grounded [36,38].

#### 4.2 Disease subtyping from omic signatures

Beyond categorical diagnosis, AI is increasingly used to resolve disease into molecular subtypes that carry prognostic and therapeutic significance. Transcriptomic and proteomic profiling of inflammatory skin reveals reproducible signatures that machine learning can exploit to stratify patients. In psoriasis, gene-expression analysis defines a dominant IL-17/IL-23 inflammatory programme; in atopic dermatitis, profiling exposes both the type 2 immune signature and barrier-differentiation defects, including in clinically uninvolved skin [51,52]. Such molecular stratification underpins the concept of disease endotypes - subgroups defined by mechanism rather than morphology - that may respond differentially to targeted therapy. Unsupervised clustering of multi-omic data can identify previously unrecognised patient subgroups, while supervised classifiers can assign a new patient to an established subtype from a limited biomarker panel, bringing the precision-medicine paradigm of oncology to inflammatory dermatoses [50].

#### 4.3 Case study: distinguishing psoriasis from eczema

The differentiation of psoriasis from atopic dermatitis (eczema) is a clinically meaningful and analytically illuminating test case, because the two conditions can overlap morphologically yet diverge sharply in their cytokine biology and optimal therapy [7,12]. Approaches operating purely on images have achieved respectable

accuracy in separating these and other inflammatory dermatoses [54]. The more powerful strategy, however, exploits their molecular separability: classifiers trained on cytokine or transcriptomic profiles can discriminate the Th17/IL-17-skewed signature of psoriasis from the Th2/IL-13-skewed signature of atopic dermatitis with high fidelity, and these biochemical features can be fused with imaging to resolve cases that are ambiguous on inspection alone [52]. This combined image-plus-biochemistry strategy exemplifies the central thesis of this review: that AI delivers its greatest diagnostic value when it integrates what the lesion looks like with what the lesion is, biochemically.

#### 4.4 Performance, validation, and the reality gap

Reported accuracies for dermatological AI are frequently impressive, with area-under-the-curve values for melanoma classification commonly exceeding 0.90 in controlled studies [34,35,53]. These figures must, however, be interpreted with caution. Much of the published performance derives from retrospective evaluation on curated datasets that may not represent the spectrum of real-world presentations, skin phototypes, or image quality [79,81]. Models can exploit spurious correlations - for instance, learning to associate surgical skin markings or rulers in dermoscopic images with malignancy rather than learning genuine biological features [83]. Prospective, multi-site validation on demographically diverse populations remains comparatively rare, and head-to-head performance often degrades on external data. The field is accordingly maturing from a narrow focus on benchmark accuracy towards an emphasis on generalisability, calibration, and the demonstration that AI augments rather than merely matches clinicians; studies showing that decision support raises the performance of practitioners above either human or machine alone point to the most realistic near-term clinical model [82].

#### 4.5 AI and non-invasive biochemical imaging

A frontier of particular relevance to biochemistry is the AI-assisted interpretation of non-invasive imaging modalities that report on tissue

composition rather than surface morphology alone. Reflectance confocal microscopy resolves cellular and architectural detail in vivo, approaching histology without biopsy; optical coherence tomography images subsurface structure; and vibrational spectroscopic techniques such as Raman and infrared spectroscopy directly probe molecular composition - lipids, proteins, and metabolites - within the skin. Electrical impedance spectroscopy, the basis of a regulator-approved melanoma device, infers cellular and biochemical tissue properties from their electrical behaviour [53,71]. Each of these modalities generates

complex, high-dimensional signals that are difficult for humans to interpret but well matched to machine learning. By learning the mapping between these biophysical and biochemical signatures and tissue diagnosis, AI can transform non-invasive imaging into a quantitative, point-of-care biochemical assay, reducing the need for invasive biopsy and enabling longitudinal monitoring. The convergence of such modalities with conventional dermoscopy and histology within multimodal models represents one of the more promising near-term routes to biochemically grounded, AI-augmented diagnosis.

**Table 2. Representative artificial-intelligence models for the diagnosis and biochemical profiling of skin diseases, with input data and reported performance.**

Skin disorder / task	AI model type	Input data (incl. biochemical features)	Reported accuracy	Reference
Melanoma (dermoscopy)	Deep CNN	Dermoscopic images; lesion morphology features	AUC > 0.90; matched/exceeded most dermatologists	Haenssle 2018 [34]; Brinker 2019 [35]
Skin cancer (general)	Deep CNN	Clinical + dermoscopic images	Dermatologist-level classification	Esteva 2017 [29]
Pigmented lesions	ML algorithms vs human readers	Dermoscopic image archive (ISIC/HAM10000)	Algorithms equalled or exceeded human readers	Tschandl 2019 [36]; 2018 [33]
Multi-class skin disease	Deep learning differential diagnosis	Clinical images + case metadata	Comparable to dermatologists across many conditions	Liu 2020 [38]; Han 2020 [82]
Psoriasis vs eczema	CNN / molecular classifier	Clinical images; cytokine and transcriptomic profiles	High discrimination using Th17 vs Th2 signatures	Wu 2020 [54]; Suarez-Farinas 2011 [52]

**5. AI in Drug Discovery and Personalised Treatment**

The same computational advances reshaping diagnosis are accelerating the discovery of new dermatological therapeutics and the personalisation of existing ones. AI compresses the timelines and reduces the attrition of drug

development by predicting molecular properties, prioritising targets and compounds, and forecasting which patient will benefit from which agent. Figures 2 and 3 depict, respectively, an AI-driven drug-repurposing pipeline and a personalised-treatment workflow built around biochemical profiling.

### 5.1 Target identification and virtual screening

Drug discovery begins with the identification and validation of a molecular target, a process that AI accelerates by integrating genomic, transcriptomic, and network data to nominate the pathways most causally implicated in disease [56]. Once a target is selected, machine learning enables high-throughput virtual screening, predicting the binding affinity, selectivity, and drug-like properties of vast compound libraries far faster and more cheaply than physical assays [56]. The transformative advance in structural biology - highly accurate computational prediction of protein structure - has further expanded the universe of tractable targets by providing reliable three-dimensional models for structure-based design [58]. Deep generative models can propose entirely novel molecular structures optimised for a desired target, an approach that has yielded potent candidate inhibitors on timescales of weeks rather than years [59], while related methods have surfaced antibacterial compounds with mechanisms distinct from existing antibiotics - directly relevant to the microbial dimension of acne and to skin and soft-tissue infection [57]. Applied to dermatology, these tools can be directed at skin-specific biochemical pathways such as tyrosinase in pigmentary disease, sebaceous lipogenic enzymes in acne, or the kinases and cytokine receptors central to inflammatory dermatoses.

Beyond identifying candidate molecules, AI increasingly addresses the attributes that most often derail development: absorption, distribution, metabolism, excretion, and toxicity. Machine-learning models trained on large pharmacological datasets can predict these ADMET properties early, filtering out compounds likely to fail on safety or pharmacokinetic grounds before costly synthesis and assay [56]. For dermatological agents intended for topical delivery, prediction of skin permeability, percutaneous absorption, and local irritancy is especially valuable, because a molecule must traverse the formidable barrier of the stratum corneum to reach its biochemical target while avoiding sensitisation. Models that estimate the

partitioning and diffusion of small molecules through skin lipids, and that flag structural alerts for contact allergy, allow medicinal chemists to optimise candidates for the cutaneous route specifically. Coupled with generative design, such predictive filtering creates a closed optimisation loop in which proposed structures are simultaneously tuned for target affinity, selectivity, synthetic accessibility, and a favourable cutaneous safety profile, concentrating experimental effort on the most promising chemical space [57,59].

### 5.2 Predicting response to biologics and targeted agents

The proliferation of biologics has made the matching of patient to therapy one of the most pressing problems in clinical dermatology. Anti-IL-17, anti-IL-23, and anti-TNF agents transform psoriasis for many patients, and anti-IL-4-receptor therapy does the same for atopic dermatitis, yet response is heterogeneous and the optimal agent for a given patient cannot reliably be predicted from clinical features alone [61,62,63,64]. Multi-omic data offer a solution. Pharmacogenomic studies have linked specific genetic variants to differential response and toxicity of biologic and conventional therapies in psoriasis [65], and longitudinal molecular profiling has shown that early changes in lesional gene expression can predict eventual clinical response to biologic agents, potentially shortening the trial-and-error interval [66]. Machine learning models that integrate baseline transcriptomic, proteomic, serological, and clinical features are being developed to forecast both efficacy and adverse events, moving the selection of targeted therapy from empiricism towards prediction. In atopic dermatitis, baseline biomarkers such as serum immunoglobulin E and filaggrin mutation status provide candidate predictive features for such models [10,12].

Prediction of harm is as important as prediction of benefit. Biologics can provoke immunogenic responses, with anti-drug antibodies that neutralise efficacy and occasionally precipitate adverse reactions, and conventional immunosuppressants carry organ-specific toxicities that vary with genotype [65]. Machine-

learning models that integrate pharmacogenomic variants, baseline immune and metabolic markers, and prior treatment history can stratify patients by their risk of secondary loss of response or serious adverse events, informing both agent selection and monitoring intensity. This dual prediction of efficacy and safety is the analytical core of precision dermatology: rather than cycling a patient through successive biologics over months while disease activity and quality of life deteriorate, a model can rank the available agents by expected net benefit for that individual, accounting for their particular biochemical phenotype. Realising this vision at scale will require prospective, multi-centre studies that link standardised baseline multi-omic profiling to longitudinal outcomes, but the conceptual framework and early evidence are already in place [62,66].

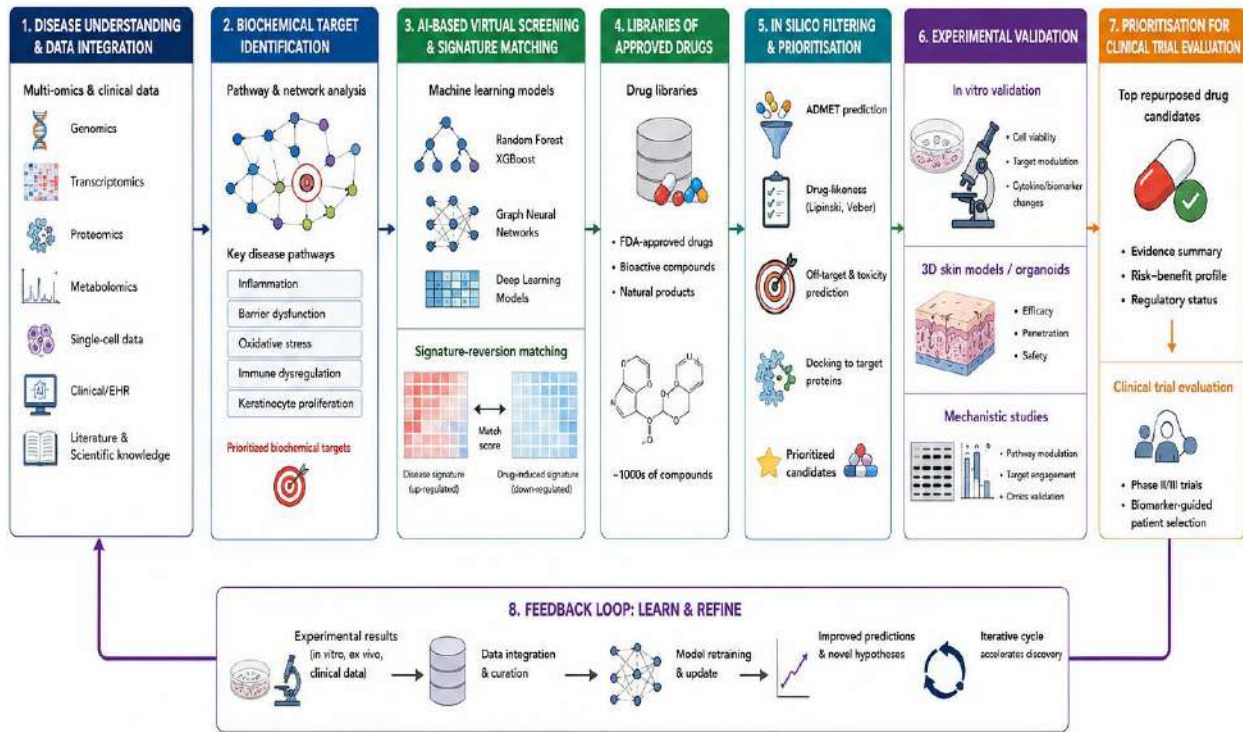
### 5.3 Drug repurposing for common and rare dermatoses

Developing a new drug de novo is slow and expensive; repurposing an approved drug to a new indication shortcuts much of the preclinical and safety pipeline [60]. AI is especially powerful here, because it can systematically compare the molecular signatures of disease with the transcriptomic effects of thousands of existing compounds and predict which drugs might reverse a pathological signature. This signature-reversion strategy, together with network-based and knowledge-graph methods that connect drugs, targets, and diseases, has generated repurposing hypotheses across medicine [56,60]. For dermatology the implications are substantial. The

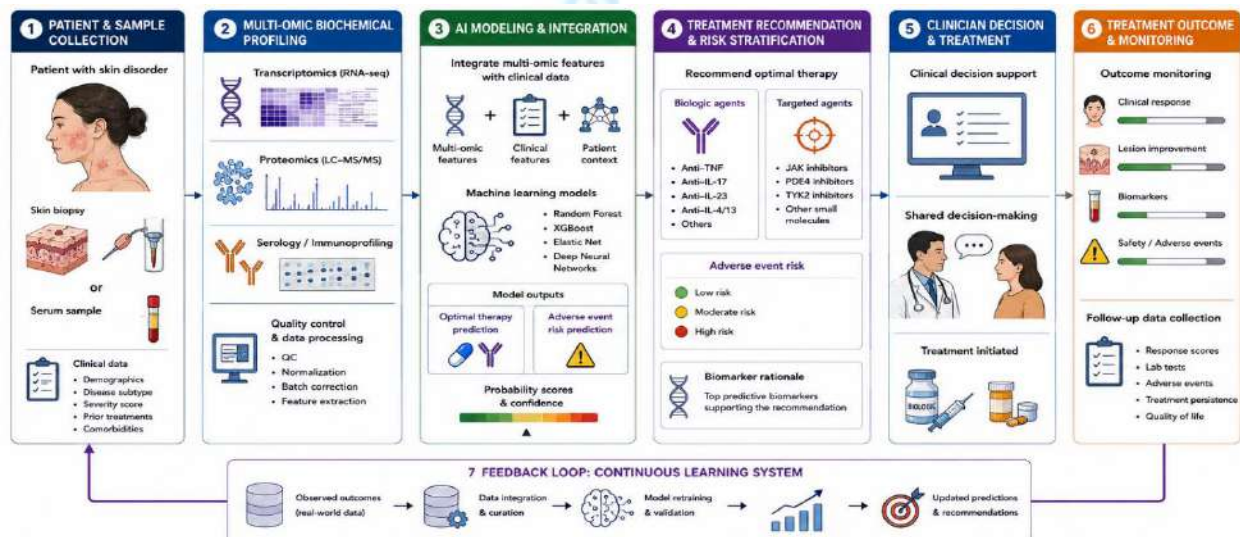
repositioning of JAK inhibitors - originally developed for haematological and rheumatological indications - to vitiligo, alopecia areata, and atopic dermatitis is a recent, clinically validated example of the underlying logic, in which a shared interferon- or cytokine-driven biochemical mechanism justifies cross-indication use [18]. For rare dermatological conditions, where dedicated drug development is economically unattractive, AI-driven repurposing may be the most realistic route to effective therapy.

### 5.4 Closing the loop: from prediction to validation

An AI-generated hypothesis, whether a novel compound or a repurposing candidate, is only a starting point. Predicted molecules must be validated in biochemical and cellular assays, in increasingly sophisticated tissue models, and ultimately in clinical trials [56]. The most credible pipelines therefore embed AI within an iterative, experimentally grounded cycle: computational prediction informs targeted experiments, whose results are fed back to refine the models, progressively improving predictive accuracy. Human skin-equivalent and organoid systems, discussed in Section 7, provide physiologically relevant platforms for this validation, and the coupling of AI prediction with such models is beginning to shorten the path from biochemical target to candidate therapy [88]. Representative trials that integrate AI-based prediction with biochemical interventions are summarised in Table 4.



**Figure 2.** AI-driven drug-repurposing pipeline for a skin disorder. Biochemical target identification from disease pathway analysis feeds machine-learning-based virtual screening and signature-reversion matching against libraries of approved drugs; predicted candidates undergo in silico filtering, in vitro and tissue-model validation, and prioritisation for clinical-trial evaluation, with experimental results returned to refine the models.



**Figure 3.** AI workflow for personalised dermatological treatment. A patient biopsy or serum sample undergoes multi-omic biochemical profiling (transcriptomics, proteomics, serology); a machine-learning model integrates these features with clinical data to predict the optimal biologic or targeted agent and likely adverse events; the clinician makes a treatment decision; and the observed treatment outcome is fed back as new training data, creating a continuously improving learning system.

## 6. AI for Monitoring and Curing: Real-World Applications

The translation of AI from research datasets into everyday care depends on tools that operate continuously, at the point of need, and within existing clinical workflows. This section surveys real-world applications in monitoring, treatment optimisation, and emerging digital therapeutics. Table 3 lists clinically validated AI tools and digital therapeutics together with the biochemical parameters they address.

### 6.1 Wearable sensors and smartphone-based monitoring

Chronic skin diseases fluctuate over time, yet conventional care captures only intermittent snapshots at clinic visits. Wearable and smartphone-based technologies promise continuous, objective monitoring of biochemically meaningful parameters [67,68]. Sensors can track skin hydration and transepidermal water loss as proxies for barrier integrity in atopic dermatitis, surface lipid and sebum levels in acne, erythema and temperature as correlates of inflammation, and scratching behaviour as an objective measure of pruritus. Emerging epidermal and microfluidic sensors extend this to the molecular level, sampling sweat and interstitial fluid for metabolites, electrolytes, and inflammatory mediators, thereby bringing genuine biochemical monitoring out of the laboratory and onto the skin [67]. AI is essential to make sense of the resulting high-frequency, multimodal data streams, detecting early signs of flare, quantifying disease activity, and personalising the timing of intervention. Smartphone cameras, combined with on-device or cloud-based AI, also enable patients to monitor lesions between visits, although systematic review has shown that the accuracy of consumer skin-assessment apps is currently variable and that regulatory and validation standards must improve before they can be relied upon for decisions such as melanoma triage [70].

The molecular frontier of this field is the most consequential for biochemistry. Whereas first-generation wearables inferred disease activity from physical surrogates such as temperature and water

loss, the newest epidermal biosensors perform genuine chemical analysis at the skin surface, using enzymatic and electrochemical detection to quantify analytes including lactate, glucose, urea, cortisol, and cytokines in sweat and interstitial fluid [67]. For inflammatory dermatoses, continuous tracking of locally released cytokines and oxidative-stress markers could allow flares to be detected and treated before they become clinically apparent; for acne, real-time measurement of surface lipid composition could guide sebum-modulating therapy. The volume, noise, and inter-individual variability of these signals make AI indispensable for extracting clinically actionable patterns, calibrating against reference assays, and distinguishing meaningful biochemical change from artefact. As these molecular wearables mature, the boundary between diagnosis, monitoring, and therapy adjustment will blur into a single continuous, biochemically informed feedback loop, realising in practice the closed-loop vision that has long been aspirational in chronic disease management [68].

### 6.2 AI-guided optimisation of phototherapy and topical regimens

Treatment of inflammatory and pigmentary dermatoses often involves regimens - phototherapy schedules, topical agent selection and tapering - that are currently individualised by clinical judgement and trial and error. AI can optimise these regimens by learning, from longitudinal data, the relationships between treatment parameters, biochemical and clinical markers of disease activity, and outcomes. In phototherapy, models can help personalise ultraviolet dosing to maximise efficacy while minimising cumulative carcinogenic and erythemogenic risk, balancing the biochemical benefits of immunomodulation against the harms of DNA damage. Reinforcement-learning frameworks are, in principle, well suited to such sequential dose-titration problems, although their clinical use awaits rigorous prospective validation [49]. Decision-support systems that recommend and adjust topical therapy on the basis of monitored barrier and inflammatory markers represent a nearer-term application, with the potential to

improve adherence and outcomes in conditions such as atopic dermatitis.

### 6.3 Digital therapeutics and validated clinical tools

A growing number of AI-based tools have achieved regulatory clearance or clinical validation, and dermatology and radiology together account for a substantial share of approved AI medical devices [71,73]. The majority of cleared dermatological tools address lesion assessment and skin-cancer risk stratification, typically operating on dermoscopic or clinical images and intended to support, rather than replace, clinician decision-making [53,71]. The regulatory landscape is evolving to accommodate the distinctive features of these products - in particular their capacity to learn and change over time, which conventional device frameworks designed for static products handle poorly [72]. Digital therapeutics that combine monitoring, behavioural support, and AI-guided treatment adjustment are an emerging category with particular promise in chronic inflammatory disease, though robust evidence of clinical benefit and biochemical effect remains to be consolidated for most.

Real-world deployment also reshapes how dermatological care is accessed. The combination of AI-based image assessment with telemedicine extends specialist-level triage to populations with limited access to dermatologists, building on an established teledermatology infrastructure that has long sought to bridge geographical gaps in specialist care [69]; this is a matter of particular consequence given the global maldistribution of the workforce and the time-critical nature of melanoma detection [81]. In such settings the value of AI lies less in marginal accuracy gains over an expert and more in providing a consistent, scalable first assessment where no expert is otherwise available. Yet these same deployments expose the field's central equity risk: tools trained predominantly on lighter skin phototypes may underperform precisely in the populations that

stand to benefit most from improved access, and uncritical reliance could widen rather than narrow disparities [79,82]. Responsible deployment therefore demands not only regulatory clearance but ongoing post-market surveillance of performance across skin types, transparent communication of a tool's validated scope, and integration designed to augment rather than substitute for clinical judgement, with clear escalation pathways for ambiguous or high-risk findings.

### 6.4 Towards cure: success stories and remission prediction

The ultimate aspiration is not merely to monitor disease but to achieve durable remission or cure. AI contributes to this goal by predicting which patients will achieve remission and under what regimen, enabling pre-emptive and optimally targeted therapy. In atopic dermatitis, the integration of baseline biochemical features - including serum immunoglobulin E concentrations, filaggrin mutation status, and cytokine profiles - with clinical data offers a route to forecasting which patients will respond durably to barrier-directed and immunomodulatory therapy [10,12,63]. In psoriasis, early molecular response profiling can predict eventual clinical clearance and guide the choice and timing of biologic therapy, helping to convert chronic management into sustained, treatment-free remission for a subset of patients [66]. In vitiligo, biomarker-guided use of JAK inhibitors illustrates how a mechanistically targeted, biochemically rationalised intervention can reverse a previously intractable disease, and AI-based prediction of repigmentation response could further refine patient selection [18]. These examples remain early, and durable cure is achieved in only a minority of patients, but they delineate a realistic trajectory in which AI-augmented, biochemically informed precision dermatology progressively improves the odds of lasting disease control.

**Table 3. Clinically validated or regulatory-cleared AI tools and digital diagnostics for skin disorders.**

Tool / device	Developer	AI technique	Parameters assessed	Indication	Regulatory status*
DERM	Skin Analytics	Deep ensemble CNN	Lesion morphology (dermoscopic/clinical images)	Skin cancer triage	CE / UKCA marked
SkinVision	SkinVision	CNN (fractal/feature analysis)	Lesion images; risk features	Skin cancer risk assessment	CE marked
Nevisense	SciBase	Electrical impedance spectroscopy + ML	Tissue bioimpedance (cellular/biophysical)	Melanoma detection	FDA approved (PMA)
Derm Assist	Google Health	Deep learning system	Clinical images + metadata	Dermatological condition triage	CE marked (EU)
MetaOptima / molemap-type platforms	Various	CNN + total-body imaging	Serial lesion imaging	Lesion monitoring	Jurisdiction-dependent

**Table 4. Representative trial designs integrating AI-based prediction with biochemical interventions in dermatology.**

Trial design (representative)*	Disease	AI intervention	Biochemical endpoint	Phase / status
AI biologic-response predictor	Psoriasis	ML model predicting response to anti-IL-17/IL-23 from baseline transcriptomics	Lesional IL-17/IL-23 gene-expression change; PASI response	Translational / II
Biomarker-guided dupilumab selection	Atopic dermatitis	Multi-omic classifier for type 2 endotype	Serum IgE, TARC/CCL17; EASI response	II / observational
AI dosing of phototherapy	Psoriasis / vitiligo	RL-based UV dose optimisation	Erythema/inflammatory markers; repigmentation	Pilot / feasibility
AI-guided JAK inhibitor use	Vitiligo	Predictive model of repigmentation response	CXCL10; melanin index	II / III
Wearable sensor + AI flare prediction	Atopic dermatitis	Time-series model on hydration/TEWL data	Transepidermal water loss; barrier markers	Feasibility / device study

7. Challenges and Future Directions

For all its promise, AI in dermatology remains far from routine clinical reality. The barriers to adoption are not principally computational but concern data, model behaviour, equity, regulation, and the irreducible complexity of biology. Figure 4 organises these barriers as a cause-and-effect (Ishikawa) diagram, and this section examines each in turn before sketching the technologies likely to shape the next decade.

### 7.1 Data heterogeneity and the absence of biochemical ontologies

AI models are only as good as the data on which they are trained, and dermatological data are notoriously heterogeneous. Imaging datasets differ in acquisition device, lighting, and resolution; omic datasets differ in platform, normalisation, and batch effects; and clinical records vary in completeness and coding. Crucially, biochemical data lack the standardised, interoperable ontologies that would allow measurements from different laboratories and platforms to be combined reliably, so that a metabolite or protein quantified in one study cannot be assumed equivalent to the same analyte in another. Missing values are pervasive in real-world biochemical datasets, and small sample sizes - especially for rare diseases and for non-European ancestries - limit statistical power and generalisability. Addressing these problems will require concerted investment in data standardisation, shared reference datasets, and harmonised biochemical ontologies, alongside privacy-preserving methods such as federated learning that allow models to be trained across institutions without centralising sensitive patient data [84,85].

### 7.2 Interpretability and the black-box problem

Many high-performing models, deep neural networks especially, are opaque: they yield predictions without an intelligible account of why. In a domain where decisions carry serious consequences, this opacity is a genuine obstacle to clinician trust, regulatory approval, and the detection of spurious reasoning [74]. Post-hoc explanation methods - feature-attribution techniques and saliency mapping that highlight the image regions or molecular features driving a

prediction - provide partial transparency and can expose models that rely on artefacts rather than biology [75,76,83]. A complementary view holds that for high-stakes decisions, inherently interpretable models should be preferred over opaque models with bolted-on explanations wherever performance permits [74]. The integration of biochemical knowledge offers a particularly attractive route to interpretability: a model that bases a psoriasis-versus-eczema decision on a recognised cytokine signature is intrinsically more explicable, and more trustworthy, than one that reasons over inscrutable pixel statistics.

### 7.3 Bias, equity, and ethical concerns

AI systems can encode and amplify the biases present in their training data, with serious equity implications. In dermatology this is acute: most large image datasets substantially under-represent darker skin phototypes, so that models trained on them may perform poorly precisely for the populations already underserved by specialist care [78,79,80]. Algorithmic bias has been demonstrated to disadvantage minority groups in other areas of healthcare, and there is every reason to expect the same in dermatology unless datasets are deliberately diversified and performance is audited across phototypes and demographic groups [78,80]. Beyond bias, the use of sensitive biochemical and genomic data raises pressing questions of consent, privacy, ownership, and the potential for discrimination, all of which demand robust governance [77]. Building equitable systems will require representative data, disaggregated performance reporting, and the meaningful involvement of affected communities in design and evaluation [79].

### 7.4 Regulatory and translational hurdles

Translating a validated algorithm into a clinically deployed tool requires navigating a regulatory framework that is still adapting to the peculiarities of AI [72,73]. Adaptive models that continue to learn after deployment challenge traditional notions of a fixed, certified device, and regulators are developing new frameworks to permit controlled model updating while preserving safety [72]. Reimbursement pathways, integration with

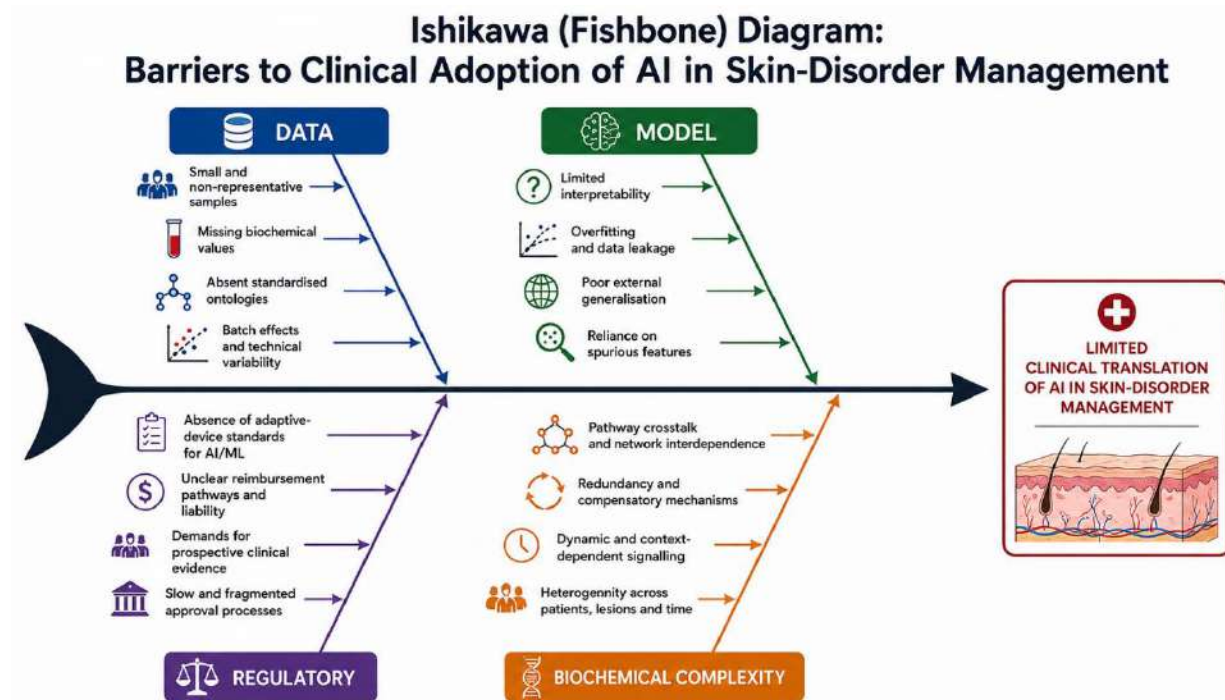
electronic health systems, liability for AI-influenced decisions, and the need for prospective evidence of clinical and economic benefit all add friction. The comparatively small number of dermatological AI tools in routine use, despite a large research literature, reflects the gap between demonstrating accuracy on a dataset and demonstrating value in a clinic [71,81].

### 7.5 Integration with emerging technologies

The future of dermatological AI will be shaped by its convergence with other advancing technologies. Organ-on-chip and human skin-equivalent systems provide physiologically realistic platforms on which AI-predicted hypotheses can be tested rapidly and ethically, compressing the validation cycle in drug discovery [86]. Skin organoids generated from pluripotent stem cells, capable of producing complex appendage-bearing tissue, offer unprecedented models for studying pathophysiology and screening therapies [88]. Genome-editing technologies, particularly CRISPR-based systems, enable both functional dissection of candidate biochemical targets through high-throughput screens and the prospect of correcting causal mutations in monogenic skin diseases [87]. The integration of AI with these wet-lab platforms - using models to design experiments and interpret their results - promises a tightly coupled, iterative discovery engine. In parallel, foundation and multimodal models that reason

jointly over images, molecular profiles, and text point towards generalist dermatological AI capable of supporting the full clinical pathway, provided the challenges of data, interpretability, equity, and regulation are met [89,90].

A further enabling development concerns the infrastructure through which models are trained on sensitive biochemical and clinical data. Privacy-preserving paradigms such as federated learning, in which models are trained across multiple institutions without the raw patient data ever leaving its source, offer a route to assembling the large, diverse datasets that robust and equitable models demand while respecting the ethical and legal constraints on genomic and health information [84,85]. Combined with techniques such as differential privacy and synthetic data generation, these approaches could allow geographically and ethnically diverse cohorts - essential for fairness across skin phototypes - to contribute to model development without centralising their data [77]. The maturation of this infrastructure is as important as advances in model architecture: the most sophisticated algorithm is of limited value if it cannot be trained on data that are representative of the populations it will serve. Building trustworthy, interoperable, and privacy-respecting data ecosystems is therefore a prerequisite for, rather than an adjunct to, the clinical realisation of AI-augmented precision dermatology.



**Figure 4.** Ishikawa (fishbone) diagram of barriers to clinical adoption of AI in skin-disorder management. Four principal categories of cause converge on the central problem of limited clinical translation: data (small and non-representative samples, missing biochemical values, absent standardised ontologies, batch effects); model (limited interpretability, overfitting, poor external generalisation, reliance on spurious features); regulatory (absence of adaptive-device standards, unclear reimbursement and liability, demands for prospective evidence); and biochemical complexity (pathway crosstalk, redundancy, dynamic and context-dependent signalling).

## 8. Conclusion and Perspective

Skin disorders are biochemically diverse, clinically burdensome, and globally prevalent, and the conventional tools used to diagnose and treat them are increasingly outmatched by the complexity of the underlying biology. Across this review a consistent picture has emerged: artificial intelligence is most powerful in dermatology not as a stand-alone image classifier but as an integrative technology that fuses imaging with the biochemical and multi-omic data defining disease mechanism. It is this fusion that allows AI to identify novel biomarkers, sharpen diagnosis beyond the limits of morphology, stratify disease into mechanistically meaningful endotypes, predict and personalise response to targeted therapy, and accelerate the discovery and repurposing of drugs.

Equally clear, however, are the conditions under which this promise will or will not be realised. AI does not transcend the quality of its data, the representativeness of its training populations, or the soundness of the biochemical knowledge it encodes. Without standardised biochemical ontologies, diverse and well-curated datasets, interpretable models, rigorous prospective validation, attention to equity across skin phototypes, and adaptive regulatory frameworks, even the most accurate algorithm will fail to translate into better care - and may entrench existing disparities. The black-box problem and algorithmic bias are not peripheral caveats but central determinants of whether AI helps or harms.

It is worth situating these developments against the historical arc of dermatological science. For most of its history, dermatology was a discipline of

the visible surface, in which diagnosis rested on the trained eye and therapy on broad, often empirical interventions. The molecular era reframed skin disease as the readout of definable biochemical pathways - cytokine cascades, enzymatic dysregulation, oncogenic mutations - and gave rise to the targeted therapies that have transformed conditions such as psoriasis and melanoma. Artificial intelligence represents the next inflection in this trajectory, not by displacing biochemistry but by rendering its vast, high-dimensional data tractable. Where a clinician can hold a handful of variables in mind, a model can integrate thousands of molecular features across genome, transcriptome, proteome, and metabolome, together with images and longitudinal clinical histories, and detect the patterns that connect mechanism to phenotype to treatment response. The synthesis of biochemistry and computation thus completes a logical progression: from describing what the skin looks like, to understanding why it behaves as it does, to predicting what it will do and how best to intervene.

Realising the benefits of AI-augmented precision dermatology will require genuine multidisciplinary collaboration. Biochemists and molecular biologists must define the pathways and biomarkers that give models their mechanistic grounding; clinicians must frame the questions that matter and steward the responsible deployment of these tools; and data scientists must build models that are accurate, interpretable, equitable, and robust. The convergence of multimodal AI with organ-on-chip systems, skin organoids, genome editing, and wearable biosensors points towards a digital dermatology ecosystem in which biochemical data acquired continuously and non-invasively are interpreted in real time to guide diagnosis, monitoring, and individualised therapy. Such a system would not replace biochemical understanding but would extend and operationalise it, turning mechanistic insight into more effective, more personalised, and more accessible care.

### Declarations

**Funding:** There is no Funding Source for this paper.

**Conflicts of interest:** the authors declare no competing interests relevant to this review

**Author contributions:** all authors contributed to the conception, drafting, and critical revision of the manuscript.

**Data availability:** not applicable to this review.

### REFERENCES

- Hay RJ, Johns NE, Williams HC, et al. The global burden of skin disease in 2010: an analysis of the prevalence and impact of skin conditions. *J Invest Dermatol.* 2014;134(6):1527-1534.
- Seth D, Cheldize K, Brown D, Freeman EE. Global burden of skin disease: inequities and innovations. *Curr Dermatol Rep.* 2017;6(3):204-210.
- Karimkhani C, Dellavalle RP, Coffeng LE, et al. Global skin disease morbidity and mortality: an update from the Global Burden of Disease Study 2013. *JAMA Dermatol.* 2017;153(5):406-412.
- Flohr C, Hay R. Putting the burden of skin diseases on the global map. *Br J Dermatol.* 2021;184(2):189-190.
- Nestle FO, Kaplan DH, Barker J. Psoriasis. *N Engl J Med.* 2009;361(5):496-509.
- Boehncke WH, Schon MP. Psoriasis. *Lancet.* 2015;386(9997):983-994.
- Lowes MA, Suarez-Farinas M, Krueger JG. Immunology of psoriasis. *Annu Rev Immunol.* 2014;32:227-255.
- Hawkes JE, Chan TC, Krueger JG. Psoriasis pathogenesis and the development of novel targeted immune therapies. *J Allergy Clin Immunol.* 2017;140(3):645-653.
- Weidinger S, Novak N. Atopic dermatitis. *Lancet.* 2016;387(10023):1109-1122.

- Palmer CN, Irvine AD, Terron-Kwiatkowski A, et al. Common loss-of-function variants of the epidermal barrier protein filaggrin are a major predisposing factor for atopic dermatitis. *Nat Genet.* 2006;38(4):441-446.
- Weidinger S, Beck LA, Bieber T, Kabashima K, Irvine AD. Atopic dermatitis. *Nat Rev Dis Primers.* 2018;4(1):1.
- Brunner PM, Guttman-Yassky E, Leung DY. The immunology of atopic dermatitis and its reversibility with broad-spectrum and targeted therapies. *J Allergy Clin Immunol.* 2017;139(4S):S65-S76.
- Williams HC, Dellavalle RP, Garner S. Acne vulgaris. *Lancet.* 2012;379(9813):361-372.
- Zaenglein AL. Acne vulgaris. *N Engl J Med.* 2018;379(14):1343-1352.
- Dreno B. What is new in the pathophysiology of acne, an overview. *J Eur Acad Dermatol Venereol.* 2017;31(Suppl 5):8-12.
- Smith RN, Mann NJ, Braue A, Makelainen H, Varigos GA. A low-glycemic-load diet improves symptoms in acne vulgaris patients: a randomized controlled trial. *Am J Clin Nutr.* 2007;86(1):107-115.
- Ezzedine K, Eleftheriadou V, Whitton M, van Geel N. Vitiligo. *Lancet.* 2015;386(9988):74-84.
- Frisoli ML, Essien K, Harris JE. Vitiligo: mechanisms of pathogenesis and treatment. *Annu Rev Immunol.* 2020;38:621-648.
- Pillaiyar T, Manickam M, Jung SH. Recent development of signaling pathways inhibitors of melanogenesis. *Cell Signal.* 2017;40:99-115.
- Passeron T, Picardo M. Melasma, a photoaging disorder. *Pigment Cell Melanoma Res.* 2018;31(4):461-465.
- Davies H, Bignell GR, Cox C, et al. Mutations of the BRAF gene in human cancer. *Nature.* 2002;417(6892):949-954.
- Schadendorf D, van Akkooi ACJ, Berking C, et al. Melanoma. *Lancet.* 2018;392(10151):971-984.
- Shain AH, Bastian BC. From melanocytes to melanomas. *Nat Rev Cancer.* 2016;16(6):345-358.
- Hodis E, Watson IR, Kryukov GV, et al. A landscape of driver mutations in melanoma. *Cell.* 2012;150(2):251-263.
- Chapman PB, Hauschild A, Robert C, et al. Improved survival with vemurafenib in melanoma with BRAF V600E mutation. *N Engl J Med.* 2011;364(26):2507-2516.
- Briganti S, Picardo M. Antioxidant activity, lipid peroxidation and skin diseases. What's new. *J Eur Acad Dermatol Venereol.* 2003;17(6):663-669.
- van Smeden J, Bouwstra JA. Stratum corneum lipids: their role for the skin barrier function in healthy subjects and atopic dermatitis patients. *Curr Probl Dermatol.* 2016;49:8-26.
- LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature.* 2015;521(7553):436-444.
- Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature.* 2017;542(7639):115-118.
- Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med.* 2019;25(1):44-56.
- Rajkumar A, Dean J, Kohane I. Machine learning in medicine. *N Engl J Med.* 2019;380(14):1347-1358.
- Litjens G, Kooi T, Bejnordi BE, et al. A survey on deep learning in medical image analysis. *Med Image Anal.* 2017;42:60-88.
- Tschandl P, Rosendahl C, Kittler H. The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Sci Data.* 2018;5:180161.
- Haenssle HA, Fink C, Schneiderbauer R, et al. Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists. *Ann Oncol.* 2018;29(8):1836-1842.

- Brinker TJ, Hekler A, Enk AH, et al. Deep learning outperformed 136 of 157 dermatologists in a head-to-head dermoscopic melanoma image classification task. *Eur J Cancer*. 2019;113:47-54.
- Tschandl P, Codella N, Akay BN, et al. Comparison of the accuracy of human readers versus machine-learning algorithms for pigmented skin lesion classification: an open, web-based, international, diagnostic study. *Lancet Oncol*. 2019;20(7):938-947.
- Han SS, Kim MS, Lim W, Park GH, Park I, Chang SE. Classification of the clinical images for benign and malignant cutaneous tumors using a deep learning algorithm. *J Invest Dermatol*. 2018;138(7):1529-1538.
- Liu Y, Jain A, Eng C, et al. A deep learning system for differential diagnosis of skin diseases. *Nat Med*. 2020;26(6):900-908.
- Breiman L. Random forests. *Mach Learn*. 2001;45(1):5-32.
- Cortes C, Vapnik V. Support-vector networks. *Mach Learn*. 1995;20(3):273-297.
- Libbrecht MW, Noble WS. Machine learning applications in genetics and genomics. *Nat Rev Genet*. 2015;16(6):321-332.
- Krizhevsky A, Sutskever I, Hinton GE. ImageNet classification with deep convolutional neural networks. *Commun ACM*. 2017;60(6):84-90.
- He K, Zhang X, Ren S, Sun J. Deep residual learning for image recognition. *Proc IEEE Conf Comput Vis Pattern Recognit*. 2016:770-778.
- Ronneberger O, Fischer P, Brox T. U-Net: convolutional networks for biomedical image segmentation. *Med Image Comput Comput Assist Interv*. 2015:234-241.
- Devlin J, Chang MW, Lee K, Toutanova K. BERT: pre-training of deep bidirectional transformers for language understanding. *Proc NAACL-HLT*. 2019:4171-4186.
- Lee J, Yoon W, Kim S, et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*. 2020;36(4):1234-1240.
- Wu S, Roberts K, Datta S, et al. Deep learning in clinical natural language processing: a methodical review. *J Am Med Inform Assoc*. 2020;27(3):457-470.
- Sutton RS, Barto AG. Reinforcement Learning: An Introduction. 2nd ed. Cambridge (MA): MIT Press; 2018.
- Komorowski M, Celi LA, Badawi O, Gordon AC, Faisal AA. The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nat Med*. 2018;24(11):1716-1720.
- Hasin Y, Seldin M, Lusis A. Multi-omics approaches to disease. *Genome Biol*. 2017;18(1):83.
- Tsoi LC, Iyer MK, Stuart PE, et al. Analysis of long non-coding RNAs highlights tissue-specific expression patterns and epigenetic profiles in normal and psoriatic skin. *Genome Biol*. 2015;16:24.
- Suarez-Farinas M, Tintle SJ, Shemer A, et al. Nonlesional atopic dermatitis skin is characterized by broad terminal differentiation defects and variable immune abnormalities. *J Allergy Clin Immunol*. 2011;127(4):954-964.
- Phillips M, Marsden H, Jaffe W, et al. Assessment of accuracy of an artificial intelligence algorithm to detect melanoma in images of skin lesions. *JAMA Netw Open*. 2019;2(10):e1913436.
- Wu H, Yin H, Chen H, et al. A deep learning, image based approach for automated diagnosis for inflammatory skin diseases. *Ann Transl Med*. 2020;8(9):581.
- Combalia M, Codella N, Rotemberg V, et al. BCN20000: dermoscopic lesions in the wild. *Sci Data*. 2024;11(1):641.
- Vamathevan J, Clark D, Czodrowski P, et al. Applications of machine learning in drug discovery and development. *Nat Rev Drug Discov*. 2019;18(6):463-477.

- Stokes JM, Yang K, Swanson K, et al. A deep learning approach to antibiotic discovery. *Cell*. 2020;180(4):688-702.
- Jumper J, Evans R, Pritzel A, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596(7873):583-589.
- Zhavoronkov A, Ivanenkov YA, Aliper A, et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nat Biotechnol*. 2019;37(9):1038-1040.
- Pushpakom S, Iorio F, Eyers PA, et al. Drug repurposing: progress, challenges and recommendations. *Nat Rev Drug Discov*. 2019;18(1):41-58.
- Langley RG, Elewski BE, Lebwohl M, et al. Secukinumab in plaque psoriasis - results of two phase 3 trials. *N Engl J Med*. 2014;371(4):326-338.
- Gordon KB, Blauvelt A, Papp KA, et al. Phase 3 trials of ixekizumab in moderate-to-severe plaque psoriasis. *N Engl J Med*. 2016;375(4):345-356.
- Simpson EL, Bieber T, Guttman-Yassky E, et al. Two phase 3 trials of dupilumab versus placebo in atopic dermatitis. *N Engl J Med*. 2016;375(24):2335-2348.
- Griffiths CEM, Armstrong AW, Gudjonsson JE, Barker JNWN. Psoriasis. *Lancet*. 2021;397(10281):1301-1315.
- Prieto-Perez R, Cabaleiro T, Dauden E, Ochoa D, Roman M, Abad-Santos F. Genetics of psoriasis and pharmacogenetics of biological drugs. *Autoimmune Dis*. 2013;2013:613086.
- Correa da Rosa J, Kim J, Tian S, Tomalin LE, Krueger JG, Suarez-Farinas M. Shrinking the psoriasis assessment gap: early gene-expression profiling accurately predicts response to biologic agents. *J Invest Dermatol*. 2017;137(2):305-312.
- Dunn J, Runge R, Snyder M. Wearables and the medical revolution. *Per Med*. 2018;15(5):429-448.
- Topol E. *The Topol Review: preparing the healthcare workforce to deliver the digital future*. London: Health Education England; 2019.
- Coates SJ, Kvedar J, Granstein RD. *Teledermatology: from historical perspective to emerging techniques*. *J Am Acad Dermatol*. 2015;72(4):563-574.
- Freeman K, Dinnes J, Chuchu N, et al. Algorithm based smartphone apps to assess risk of skin cancer in adults: systematic review of diagnostic accuracy studies. *BMJ*. 2020;368:m127.
- Benjamins S, Dhunoo P, Mesko B. The state of artificial intelligence-based FDA-approved medical devices and algorithms: an online database. *NPJ Digit Med*. 2020;3:118.
- US Food and Drug Administration. *Proposed regulatory framework for modifications to artificial intelligence/machine learning-based software as a medical device*. Silver Spring (MD): FDA; 2019.
- Muehlematter UJ, Daniore P, Vokinger KN. Approval of artificial intelligence and machine learning-based medical devices in the USA and Europe (2015-20): a comparative analysis. *Lancet Digit Health*. 2021;3(3):e195-e203.
- Rudin C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat Mach Intell*. 2019;1(5):206-215.
- Lundberg SM, Lee SI. A unified approach to interpreting model predictions. *Adv Neural Inf Process Syst*. 2017;30:4765-4774.
- Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-CAM: visual explanations from deep networks via gradient-based localization. *Proc IEEE Int Conf Comput Vis*. 2017:618-626.
- Char DS, Shah NH, Magnus D. Implementing machine learning in health care - addressing ethical challenges. *N Engl J Med*. 2018;378(11):981-983.

- Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366(6464):447-453.
- Daneshjou R, Smith MP, Sun MD, Rotemberg V, Zou J. Lack of transparency and potential bias in artificial intelligence data sets and algorithms: a scoping review. *JAMA Dermatol*. 2021;157(11):1362-1369.
- Adamson AS, Smith A. Machine learning and health care disparities in dermatology. *JAMA Dermatol*. 2018;154(11):1247-1248.
- Navarrete-Dechent C, Dusza SW, Liopyris K, Marghoob AA, Halpern AC, Marchetti MA. Automated dermatological diagnosis: hype or reality? *J Invest Dermatol*. 2018;138(10):2277-2279.
- Han SS, Park I, Chang SE, et al. Augmented intelligence dermatology: deep neural networks empower medical professionals in diagnosing skin cancer and predicting treatment options for 134 skin disorders. *J Invest Dermatol*. 2020;140(9):1753-1761.
- Winkler JK, Fink C, Toberer F, et al. Association between surgical skin markings in dermoscopic images and diagnostic performance of a deep learning convolutional neural network for melanoma recognition. *JAMA Dermatol*. 2019;155(10):1135-1141.
- Rieke N, Hancox J, Li W, et al. The future of digital health with federated learning. *NPJ Digit Med*. 2020;3:119.
- Sheller MJ, Edwards B, Reina GA, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Sci Rep*. 2020;10(1):12598.
- Esch MB, Bahinski A, Huh D. Organs-on-chips at the frontiers of drug discovery. *Nat Rev Drug Discov*. 2015;14(4):248-260.
- Doudna JA, Charpentier E. The new frontier of genome engineering with CRISPR-Cas9. *Science*. 2014;346(6213):1258096.
- Lee J, Rabbani CC, Gao H, et al. Hair-bearing human skin generated entirely from pluripotent stem cells. *Nature*. 2020;582(7812):399-404.
- Moor M, Banerjee O, Abad ZSH, et al. Foundation models for generalist medical artificial intelligence. *Nature*. 2023;616(7956):259-265.
- Acosta JN, Falcone GJ, Rajpurkar P, Topol EJ. Multimodal biomedical AI. *Nat Med*. 2022;28(9):1773-1784.