PUSHING THE BOUNDARIES OF COMPUTER VISION IN CLINICAL SKIN ANALYSIS

A clinical - technical analysis of our new breakthrough Skin Intelligence Engine

Pushing the Boundaries of Computer Vision in Skin Analysis

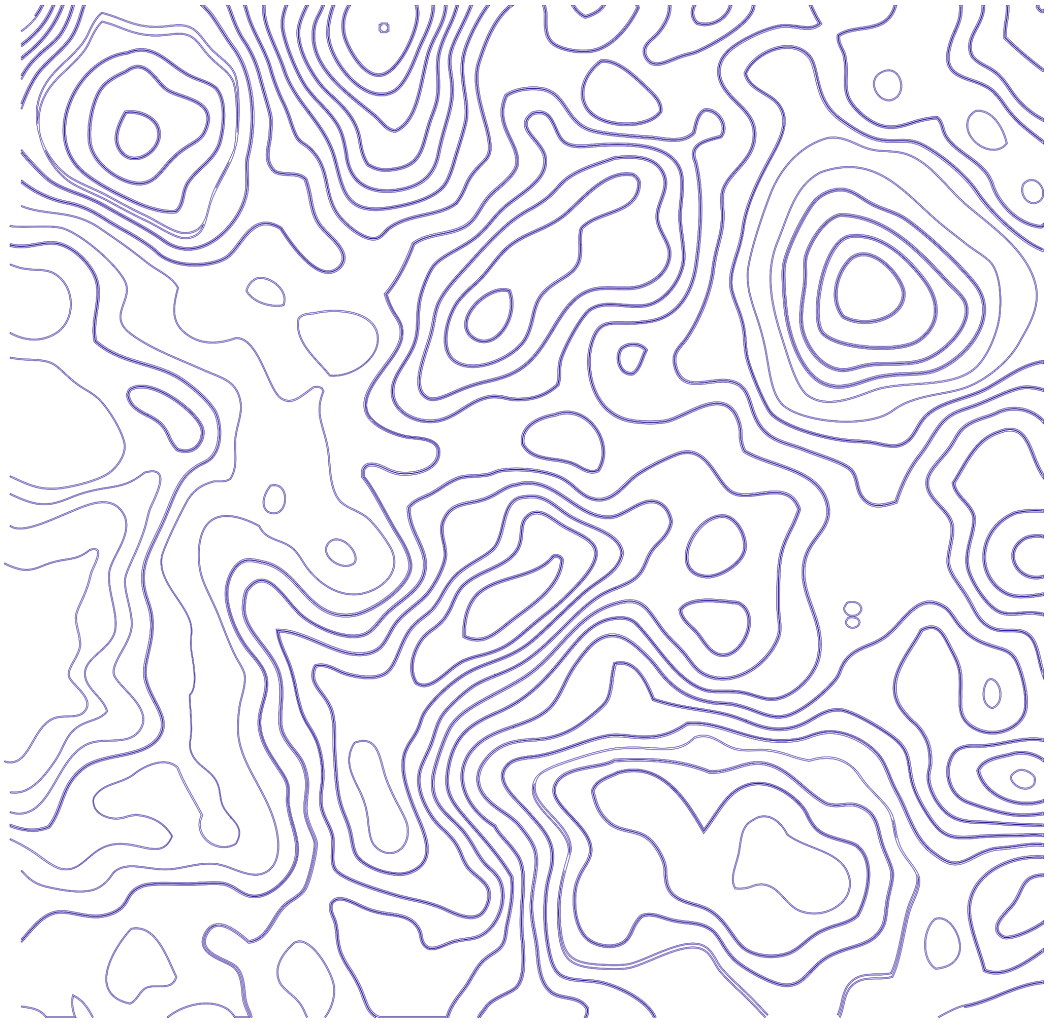

A Technical & Clinical Analysis of the Oyster Skin Intelligence Engine

Authors: Jude Jr. Chikezie, Dr Chelsea Emeh MD, Emeka Adirije, Barakat Salam
Contributors: Udoagwa Franklin EIE, Abolade Qudus IIT, Ugbala Nicholas
Subject: *Hierarchical Vision Transformers, Deterministic Pharmacological Logic, Optical Physics of Melanin-Rich Skin, Regional Demand Sensing*

Abstract

The evolution of dermatological diagnostics has reached a critical juncture where the integration of advanced computer vision and physiological modeling is no longer a luxury but a necessity for equitable global healthcare.

Automated skin analysis systems have historically been constrained by a systemic crisis of representativity that undermines their clinical utility for the majority of the world's population. Foundational datasets, such as the *International Skin Imaging Collaboration (ISIC) and HAM10000*, exhibit a profound skew toward lighter skin tones, with Fitzpatrick types I through III representing over 82% of available imagery, while types V and VI comprise less than 8%. This imbalance is not merely a statistical anomaly but a structural failure that results in "chromophore masking," a phenomenon where models trained on lighter skin mistakenly prioritize erythema as the primary indicator of inflammation. In melanin-rich skin, the high density of eumelanin acts as a broadband absorber of visible light, attenuating these vascular signals and leading to the systematic misclassification of conditions such as post-inflammatory hyperpigmentation (PIH), melasma, and early-stage eczema.

This report details the engineering of the *Oyster Skin Intelligence Engine* (SIE), a multi-task vision system utilizing a hierarchical Swin Transformer backbone to achieve clinical-grade precision across the full melanin spectrum. The SIE quantifies 33 distinct skin health metrics in under two seconds from a standard smartphone capture, incorporates a Swin-Unet segmentation architecture achieving a Dice Similarity Coefficient (DSC) of 0.9753, and deploys a deterministic pharmacological constraint solver indexing over 18,500 SKUs with more than 500 safety rules. A Lighting and Image Quality Assessment (LIQA) module increases clinically usable smartphone captures from 32% to 87%, while a regional inventory intelligence model leverages aggregate, anonymised skin telemetry to drive a 66% scan-to-cart conversion rate and a 1.6x conversion lift for enterprise partners.

1. The Socio-Technical Crisis: Chromophore Masking and Dataset Bias

The primary obstacle to equitable skin analysis is rooted in the optical physics of the epidermis and the historical bias of the data used to train diagnostic algorithms.

For decades, the medical and beauty technology industries have operated under the assumption that skin analysis tools are universally applicable, yet an examination of the underlying training data reveals that darker skin tones have been treated as edge cases. This neglect produces a cascade of clinical inaccuracies that propagate through every downstream model trained on these skewed repositories.

1.1 Quantitative Imbalance in Repositories

Systematic reviews of major dermatological repositories demonstrate an overwhelming bias toward Caucasian skin types. Analysis of the ISIC 2018 dataset reveals that the lack of representation for Fitzpatrick types V and VI is so severe that some researchers identified only a single image for type VI across entire archives. This data scarcity creates a vacuum in which models fail to learn the visual grammar of melanin-rich skin.

Dataset	Fitzpatrick I–III	Fitzpatrick IV–VI	Total Samples	Source
ISIC Archive (General)	82.06%	7.94%	17,700+	[2]
Fitzpatrick17k	75.00%	25.00%	16,577	[1]
HAM10000	95.00%	5.00%	10,015	[3]
PH2 Dataset	~100%	0.00%	200	[1]
DermNet NZ	~84.1%	~15.9%	Varies	[1]

Table 1: Fitzpatrick skin type representation across major public dermatological datasets.

The clinical cost of this misrepresentation is significant. A model that misclassifies PIH as melasma will recommend depigmentation agents when anti-inflammatory care is required, potentially exacerbating the condition. Furthermore, a failure to detect barrier

compromise on darker skin results in recommendations for exfoliating actives that further degrade the stratum corneum. These are clinical failures that erode patient trust and reinforce the perception that technology-driven skincare is fundamentally exclusionary.

1.2 The Optical Physics of Chromophore Masking

The phenomenon of chromophore masking is described by the interaction of light with the various layers of the skin. In lighter skin, the epidermis is relatively transparent, allowing the red wavelengths associated with hemoglobin in the dermal vasculature to be reflected and detected as erythema. In darker skin, the dense concentration of eumelanin in the epidermis absorbs these wavelengths before they can return to the sensor. To quantify these interactions, the SIE utilises the modified Beer-Lambert Law (mBLL), which relates the detected optical intensity to the incident intensity and the absorption and scattering properties of the tissue.

The mathematical expression for the detected intensity in a scattering medium is:

$$I = I_0 \cdot e^{\left(- \cdot \mu_a \cdot r - G\right)}$$

In this formulation, **DPF** represents the average differential pathlength factor, which accounts for the increased distance photons travel due to scattering; **μ_a** is the absorption coefficient of the chromophores (melanin and hemoglobin); **r** is the distance between the source and the detector; and **G** is a factor that accounts for geometry and scattering losses. On melanin-rich skin, the absorption coefficient μ_a is dominated by the melanin term, which effectively "masks" the smaller Δμ_a signals produced by vascular activity. To overcome this, the SIE must prioritize morphological and textural cues over simple colour-based signals.

1.3 The “Africa First” Data Strategy

Oyster’s response to this crisis was the implementation of a “Melanin-First” data strategy. Rather than attempting to retrofit diversity into an existing model by augmenting a light-skin dataset with a small number of dark-skin images, the foundational training objective was centred on Fitzpatrick types IV through VI.

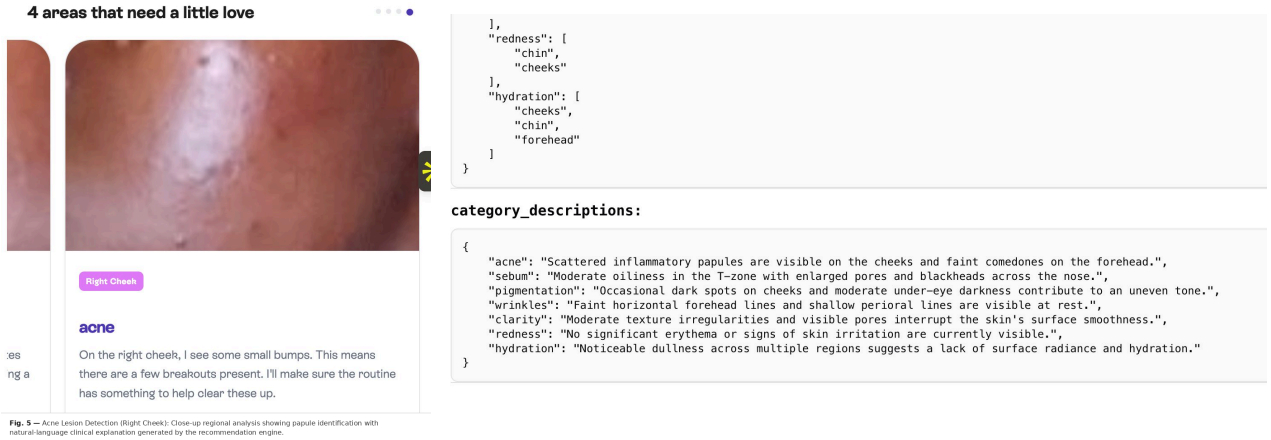

Fig. 5 — Acne Lesion Detection (Right Cheek): Close-up regional analysis showing papule identification with natural-language clinical explanation generated by the recommendation engine.

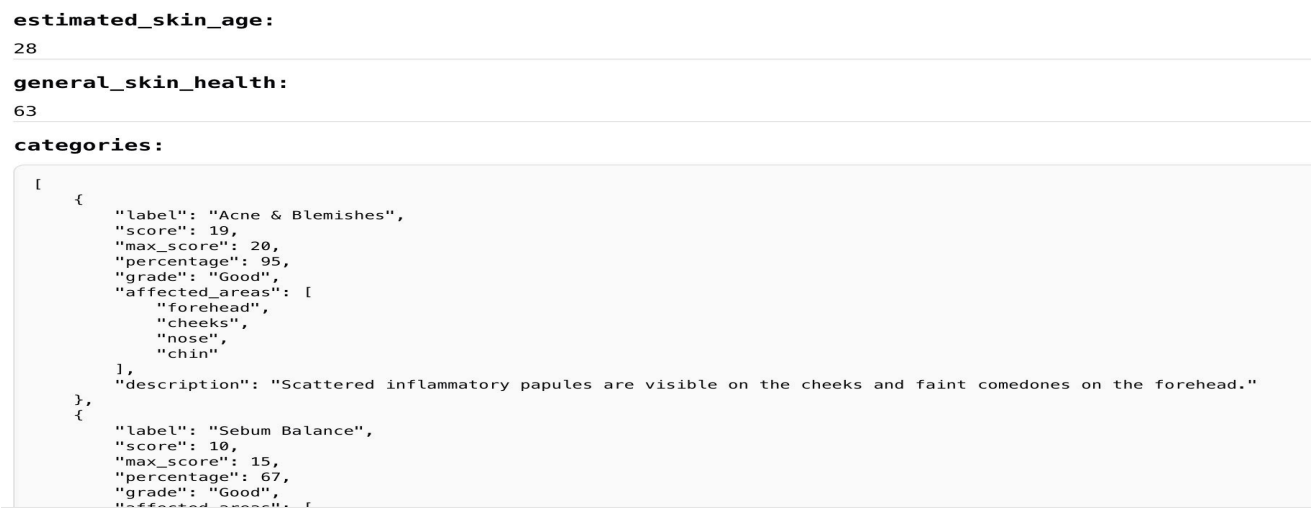

Fig. 7 — Structured Analysis Output (JSON): Raw AI engine output showing category-level scoring (0–20 scale), severity grading, region-specific observations, and category descriptions with clinical terminology.

By establishing the primary clinical hub in Johannesburg and Lagos, the SIE learns the optical complexity of melanin-rich skin as its primary visual vocabulary. This approach ensures that the model internalises the subtle changes in surface topology, specular reflectance variations, and tonal gradients that indicate skin health in populations that have been historically treated as outliers.

2. Neural Network Backbone: Hierarchical Swin Transformers

To process the complex spatial and textural features of human skin with clinical-grade accuracy, the SIE employs the **Swin Transformer** (Shifted Window Transformer) as its primary vision backbone. The requirement for high-resolution analysis (12MP+ smartphone captures) renders traditional Vision Transformers (ViTs) impractical due to their quadratic computational complexity. Standard ViTs compute global self-attention across all image patches, which for a high-resolution skin capture, would exceed any practical compute budget for mobile real-time inference.

2.1 Shifted Window Multi-head Self-Attention (SW-MSA)

The Swin Transformer introduces a hierarchical structure that limits attention computation to local, non-overlapping windows while enabling cross-window information exchange through a shifting mechanism. The self-attention within each window is defined as:

$$Attention(Q,\ K,\ V)\ =\ SoftMax\left(\frac{QK^T}{\sqrt{d_k}}\ +\ B\right)V$$

Where **Q**, **K**, and **V** are the query, key, and value matrices; **d_k** is the query/key dimension; and **B** is the learned relative position bias. The relative position bias is an essential component for skin analysis as it allows the model to learn that the spatial relationship between textural irregularities and facial landmarks (such as the nasolabial fold or the orbital rim) changes their diagnostic significance.

The computational complexity of the Swin Transformer's window-based approach is linear with respect to the image size, in contrast to the quadratic complexity of global MSA. The complexity for global MSA versus window-based MSA (W-MSA) is expressed as:

$$\Omega(MSA)\ =\ 4hwC^2\ +\ 2(hw)^2C$$

$$\Omega(W\text{-}MSA) = 4hwC^2 + 2M^2hwC$$

Where **h** and **w** are the height and width of the image in patches, **C** is the feature dimension, and **M** is the window size. By maintaining a fixed window size (e.g., M=7), the Swin Transformer achieves high efficiency, enabling the SIE to quantify 33 distinct skin health metrics in under two seconds on a standard mobile device.

2.2 Hierarchical Feature Depth and Patch Merging

The architecture builds a multi-scale representation through **patch merging**, which progressively reduces spatial resolution while increasing feature depth. This allows the SIE to mirror the diagnostic process of a clinician: initial macro-level assessment of condition distribution (Stage 4) followed by micro-level textural investigation of individual lesions (Stage 1). The hierarchy consists of four distinct stages:

1. **Stage 1:** The input image is divided into 4×4 pixel patches, each projected into a C-dimensional embedding. For a 384×384 input, this results in 96×96 = 9,216 tokens.
2. **Stage 2:** Adjacent 2×2 patches are merged, halving spatial resolution to 48×48 while doubling feature depth to 2C.
3. **Stage 3:** Further merging results in 24×24 tokens at 4C feature depth.
4. **Stage 4:** The final representation is at 12×12 tokens with 8C features, encoding the most abstract semantic information required for condition classification.

This multi-scale feature pyramid is consumed by task-specific heads for simultaneous segmentation, classification, and severity scoring.

3. Boundary Perception: Swin-Unet for Precise Segmentation

Accurate segmentation of skin lesions is the cornerstone of effective automated diagnostics. Severity scoring and subsequent treatment recommendations are directly tied to the spatial extent of a condition; thus, imprecise segmentation masks propagate errors throughout the entire inference pipeline. The SIE incorporates a **Swin-Unet** architecture, combining the hierarchical self-attention of Swin Transformers with the encoder-decoder symmetry of a U-Net.

The Swin-Unet encoder extracts multi-scale features through the patch merging hierarchy, while skip connections transfer fine-grained spatial information to the corresponding decoder stages. This preserves the boundary details that are often lost during downsampling. The architecture achieves a **Dice Similarity Coefficient (DSC) of 0.9753** in clinical benchmarks, a significant improvement over traditional CNN-based models (DeepLab V3+: DSC 0.9412; U-Net with ResNet-50 backbone: DSC 0.9589).

3.1 The Dice Similarity Coefficient (DSC)

The DSC is the primary metric for evaluating the overlap between the predicted segmentation and the clinical ground truth. It is formulated as:

$$DSC = \frac{2\,|G \cap S|}{|G| + |S|}$$

Where **G** represents the ground truth set of pixels and **S** represents the set of pixels predicted by the model. A DSC of 0.9753 indicates near-perfect alignment, which is critical for identifying sub-visible lesion margins on dark skin.

3.2 Lesion-Area Focus (LAF) Module

On melanin-rich skin, the contrast between a hyperpigmented lesion and the surrounding tissue can be minimal, making boundary detection exceptionally difficult.

The SIE addresses this through a **Lesion-Area Focus (LAF)** module located in the decoder pathway. The LAF module applies learned attention weights to boundary regions, amplifying subtle gradients at the lesion's edge. During training, the system employs boundary-weighted loss functions that penalise misclassification at the edges more heavily than in the interior, forcing the model to develop specialised edge-detection capabilities that generalise across all skin tones.

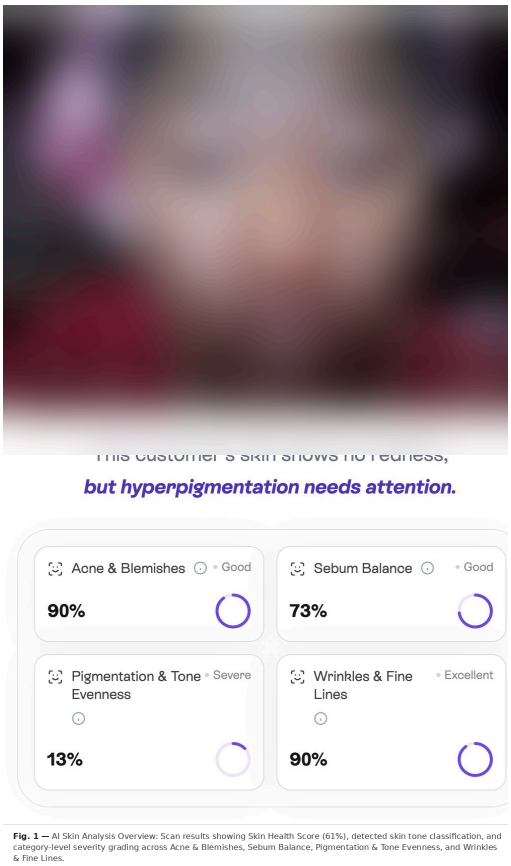

Fig. 1 — AI Skin Analysis Overview: Scan results showing Skin Health Score (61%), detected skin tone classification, and category-level severity grading across Acne & Blemishes, Sebum Balance, Pigmentation & Tone Evenness, and Wrinkles & Fine Lines.

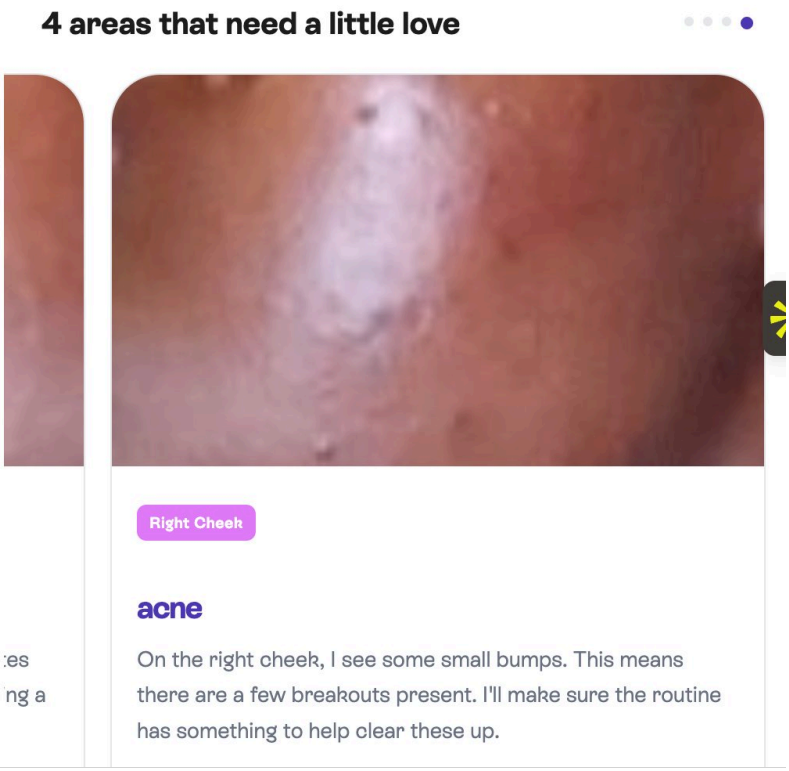

Fig. 5 — Acne Lesion Detection (Right Cheek): Close-up regional analysis showing papule identification with natural-language clinical explanation generated by the recommendation engine.

4. Tone-Adaptive Inference and ITA Calibration

To achieve tone-invariant precision, the SIE implements a tone-adaptive inference pipeline that conditions its analysis on an objective measurement of the user's skin tone. This is not a categorical demographic classification but a physical calibration step that adjusts the model's perceptual parameters based on the optics of melanin.

4.1 Individual Typology Angle (ITA) Calculation

The engine calculates the **Individual Typology Angle (ITA)** within the CIE L*a*b* colour space, using the lightness (L*) and the yellow-blue (b*) axes. The ITA is defined as:

$$ITA = arctan\left(\frac{L^* - 50}{b^*}\right) \times \left(\frac{180}{\pi}\right)$$

This metric provides a continuous, device-independent measure of constitutive pigmentation that is more granular than the traditional Fitzpatrick scale. The ITA values are mapped to qualitative skin-tone categories that inform the model's inference mode:

ITA Range (Degrees)	Fitzpatrick Category	Skin Descriptor
ITA > 55°	I	Very light
41° < ITA ≤ 55°	II	Light
28° < ITA ≤ 41°	III	Intermediate
10° < ITA ≤ 28°	IV	Tan
−30° < ITA ≤ 10°	V	Brown
ITA ≤ −30°	VI	Dark

Table 2: ITA-to-Fitzpatrick mapping used for tone-adaptive inference conditioning.

Once the ITA is established, the model adjusts its feature weights and severity thresholds. On darker skin (low ITA), the engine de-emphasises hemoglobin-based features, which are naturally attenuated by melanin, and increases the weight

assigned to **Specular Reflectance Intensity**. This enables the detection of barrier compromise and hydration levels independent of the skin's constitutive pigmentation.

4.2 Multi-Perceptual Mode Architecture

The result is a single model operating in multiple perceptual modes, each conditioned on the skin it is analysing. This is architecturally distinct from training separate models for different skin types (which would require impractically large and balanced datasets for each category) and from ignoring skin tone entirely (which produces the systematic bias documented in Section 1). One model. Multiple perceptual modes. Conditioned on the optical physics of the specific skin it is evaluating.

5. Multi-Task Learning and Uncertainty Weighting

The SIE is a **multi-task learning (MTL)** framework that simultaneously handles classification, segmentation, and biomarker quantification. Balancing these tasks is computationally challenging because they often involve different units and scales of error. If task weights are improperly balanced, a single dominant task (such as large-area segmentation) can overwhelm the training of more subtle tasks (such as micro-wrinkle detection).

5.1 Kendall Uncertainty-Weighted Loss

To address this, the SIE employs the homoscedastic uncertainty-weighting approach, which allows the model to dynamically learn the optimal weights for each task's loss function during training. The multi-task likelihood is defined by considering the observation noise (σ) for each task. The total loss for a combination of regression and classification tasks is formulated as:

$$L_{total}(W, \sigma_1, \sigma_2) = \frac{1}{2\sigma_1^2} L_{seg}(W) + \frac{1}{\sigma_2^2} L_{class}(W) + log\,\sigma_1 + log\,\sigma_2$$

Where **L_seg** and **L_class** are the individual losses for segmentation and classification, and σ_1, σ_2 are the learned noise parameters. By minimising this objective, the model effectively "attenuates" the loss of tasks with high uncertainty, preventing them from destabilising the shared feature representation. This enables the SIE to maintain high precision across all 33 skin health dimensions simultaneously.


```
    ],
    "redness": [
        "chin",
        "cheeks"
    ],
    "hydration": [
        "cheeks",
        "chin",
        "forehead"
    ]
}
```


category_descriptions:


```
{
    "acne": "Scattered inflammatory papules are visible on the cheeks and faint comedones on the forehead.",
    "sebum": "Moderate oiliness in the T-zone with enlarged pores and blackheads across the nose.",
    "pigmentation": "Occasional dark spots on cheeks and moderate under-eye darkness contribute to an uneven tone.",
    "wrinkles": "Faint horizontal forehead lines and shallow perioral lines are visible at rest.",
```

```
    "clarity": "Moderate texture irregularities and visible pores interrupt the skin's surface smoothness.",
    "redness": "No significant erythema or signs of skin irritation are currently visible.",
    "hydration": "Noticeable dullness across multiple regions suggests a lack of surface radiance and hydration."
}
```


estimated_skin_age:

28

general_skin_health:

63

categories:


```
[
    {
        "label": "Acne & Blemishes",
        "score": 19,
        "max_score": 20,
        "percentage": 95,
        "grade": "Good",
        "affected_areas": [
            "forehead",
            "cheeks",
            "nose",
            "chin"
        ],
        "description": "Scattered inflammatory papules are visible on the cheeks and faint comedones on the forehead."
    },
    {
        "label": "Sebum Balance",
        "score": 10,
        "max_score": 15,
        "percentage": 67,
        "grade": "Good",
        "affected_areas": [
```


Fig. 7 — Structured Analysis Output (JSON): Raw AI engine output showing category-level scoring (0–20 scale), severity grading, region-specific observations, and category descriptions with clinical terminology.

6. Biomarker Quantification: 33 Dimensions of Skin Health

The expansion from an initial set of 11 parameters to a **33-dimensional** skin health quantification system provides the necessary data density for truly personalised skincare. Each biomarker is extracted by a dedicated neural head operating on the shared Swin Transformer feature backbone, providing both a quantified value and a calibrated confidence score.

Biomarker	Feature / Method	Clinical Significance
Melanin Map	Subsurface Brown Mode	Distinguishes melasma from PIH via distribution depth.
Hemoglobin Map	Red Area Patterning	Detects active inflammation and rosacea patterns.
Porphyrin Count	UV-Fluorescence Mode	Measures P. acnes colonisation and sebum oxidation.
Texture Index	Specular Variance	Quantifies barrier compromise and epidermal scaling.
Hydration Index	GVR (Visible Light)	Estimates moisture content in the stratum corneum.
Micro-Wrinkles	Topographical Analysis	Detects fine lines below the human visual threshold.
TEWL Estimation	Multi-Spectral Variance	Indicates sub-visual barrier dysfunction.
Sebaceous Filament	Purity Analysis	Quantifies oxidation of sebum in the follicular opening.

Table 3: Selected biomarkers from the SIE's 33-dimensional skin health quantification system.

6.1 Grayscale Value Ratio (GVR) for Hydration Mapping

The Grayscale Value Ratio (GVR), used for hydration mapping, normalises local intensity against global facial luminance to provide a moisture index that is robust against variations in baseline skin tone and facial geometry:

$$GVR = \frac{average\ grayscale\ value\ of\ ROI}{sum\ grayscale\ value\ of\ whole\ image}$$

This normalisation ensures that the hydration index reflects actual moisture content rather than constitutive pigmentation, enabling accurate barrier integrity assessment across all skin tones. The inclusion of TEWL estimation enables the detection of barrier compromise that is not yet visible to the naked eye but that will, without intervention, progress to visible dryness, sensitivity, and increased susceptibility to irritant contact dermatitis.

7. Clinical Nuance and the Differential Diagnosis Challenge

A hallmark of the SIE is its ability to resolve visual ambiguities that challenge even board-certified dermatologists, particularly in cases where conditions present similarly on dark skin. Differential diagnosis requires the integration of morphological, spatial, and semantic cues.

7.1 Melasma vs. Post-Inflammatory Hyperpigmentation (PIH)

Melasma and PIH are frequently confused because they both present as hyperpigmented patches. However, melasma is hormonally driven and typically exhibits a symmetric, malar, or centrofacial distribution. PIH corresponds directly to the site of a prior inflammatory event (such as an acne lesion or an insect bite) and exhibits characteristic edge irregularity. The SIE resolves this by analysing the subsurface melanin distribution map; melasma tends to show homogeneous distribution across variable epidermal and dermal depths, whereas PIH deposits are spatially correlated with inflammatory markers.

7.2 Acne Vulgaris vs. Allergic Contact Dermatitis

In Fitzpatrick V–VI skin, both acne and contact dermatitis present as papules without the characteristic erythema seen in lighter skin. The SIE utilises a hierarchical classification approach:

- **Morphological Analysis:** High-resolution Stage 1 features detect the presence of oxidised keratin plugs in open comedones, a definitive signal for acne.
- **Spatial Distribution:** Stage 4 semantic features analyse the distribution pattern; acne typically follows the anatomical density of sebaceous glands in the T-zone and cheeks, while contact dermatitis presents in asymmetric, exogenous patterns that disregard anatomical boundaries.
- **Bayesian Uncertainty Quantification:** When the posterior probability of the leading differential falls below **94.2%**, the system triggers a Human-in-the-Loop review, ensuring that users receive clinical-grade guidance rather than a low-confidence automated guess. This threshold was empirically determined through ROC analysis on a held-out validation set spanning the full Fitzpatrick spectrum.

8. The Recommendation Engine: Deterministic Pharmacological Logic

The Oyster Recommendation Engine (RE) acts as a deterministic constraint solver that translates the skin assessment into a clinically coherent routine. Unlike collaborative filtering systems, which rely on user popularity and can propagate unsafe product combinations, the RE constructs routines based on pharmacological rules and ingredient compatibility.

8.1 Constraint Satisfaction Problem (CSP) Formulation

17 concern(s) detected in scans but never assigned to any product:
roughness pih_spot_count comedone nasolabial_folds shine_level inflammatory_lesions pih irregular texture pih_spot freckles depigmentation patches nasolabial lines discoloration patch perioral_lines mild_shine patches crow's feet

Concern	Times in Routines	Avg Fit %	Lowest Fit	Top Products
hyperpigmentation	4363	76.8%	34%	C Brightening Serum 30ml (461) Formulas Niacinamide Serum (228) Chemist Vivid C Collagen Boosting Radiance Serum 30ml (175)
Post-inflammatory hyperpigmentation	3583	96.4%	67%	Vita B3 Mist Toner 200ml (81) Chemist Vivid C Collagen Boosting Radiance Serum 30ml (67) Sunscreen Multi-Defense Tinted SPF 100 PA+++ | 100g (47)
T-zone shine	3157	96.8%	60%	Salicylic Acid Daily Gentle Cleanser 150ml (62) -Oil Control Moisturising Gel Cream 1.75 oz 52ml (52) -Salicylic Acid Purifying Gel Cleanser 4.0 oz 120ml (36)
comedones	3118	76.2%	28%	Salicylic Acid Daily Gentle Cleanser 150ml (473) Active Clean Gel Cleanser (194) Salicylic Acid Daily Gentle Cleanser (168)
Hyperpigmentation	2244	96.3%	67%	Vita B3 Mist Toner 200ml (54) Chemist Vivid C Collagen Boosting Radiance Serum 30ml (43) Sunscreen Multi-Defense Tinted SPF 100 PA+++ | 100g (28)
Mild dullness	2082	95.8%	66%	Chemist Vivid C Collagen Boosting Radiance Serum 30ml (32) Vita B3 Mist Toner 200ml (28) Sunscreen Multi-Defense Tinted SPF 100 PA+++ | 100g (24)
enlarged pores	2007	76.8%	35%	Active Salicylic Acid And Zinc Pore Minimizing Serum 30ml (331) 2% BHA Clarifying Exfoliant Toner (95) Eucerin Dermo Purifier Oil Control Skin Renewal Treatment (76)
Enlarged pores	1845	96.0%	64%	Zero Pore One day Cream 50ml (46) AHA/BHA Clarifying Treatment Toner (38) -Salicylic Acid (restore + clarify) facial serum 30ml (36)
under-eye darkness	1675	74.6%	34%	-Eye Repair Cream | Eye Cream for Dark Circles & Puffiness 14ml (240) FORMULAS -Ceramides Protecting Eye Cream 20ml (211) Multi-Peptide Eye Cream (157)

Fig. 8 — Concern-to-Product Mapping Engine: Recommendation accuracy metrics showing average fit percentage and top-ranked products per detected concern, with unmapped concern gap analysis (17 concerns identified).

The recommendation task is formalised as a **Constraint Satisfaction Problem (CSP)**, defined by a triple (X, D, C):

- **X:** A set of variables representing the steps in a skincare routine (e.g., Cleanser, Serum, Moisturiser, SPF).

- **D:** The domains for each variable, consisting of the available SKUs in a partner's catalogue (over **18,500 products** indexed).
- **C:** A set of **over 500 pharmacological constraints** that restrict the possible combinations of products.

The solver employs **constraint propagation** to prune invalid combinations from the search space. For instance, if Step 2 is assigned a high-strength L-ascorbic acid serum, the domain for Step 3 is immediately pruned of any products containing copper peptides, as the acidic environment of the Vitamin C serum would destabilise the peptide structure.

8.2 Multi-Layer Computational Pipeline

The recommendation logic operates across three layers:

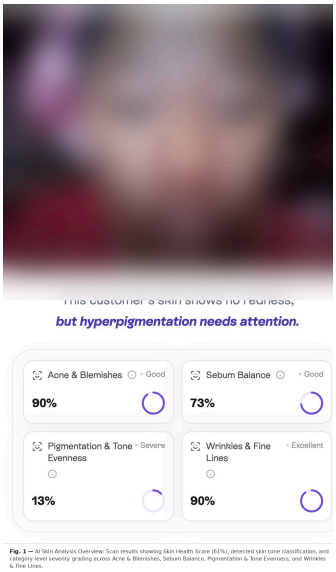

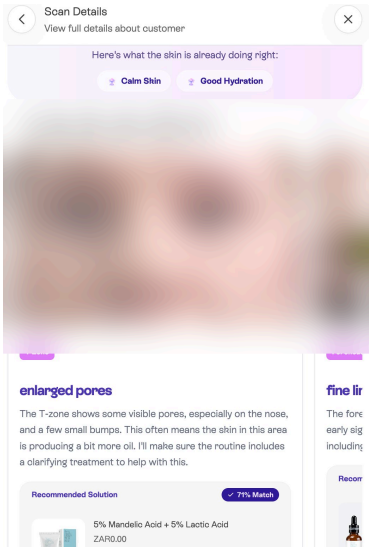

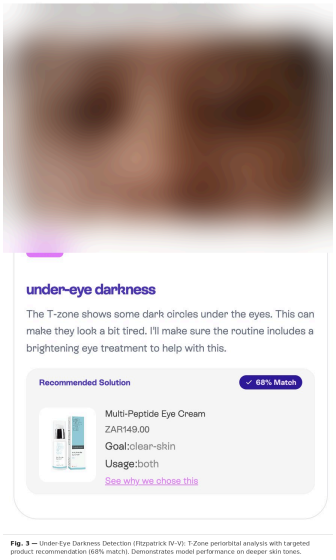

1. **Condition-to-Ingredient Mapping.** Maps detected skin metrics to specific active ingredients. On dark skin, the engine prioritises tranexamic acid and niacinamide over hydroquinone to minimise the risk of irritation-induced rebound

hyperpigmentation. Research demonstrates that while hydroquinone remains the gold standard depigmentation agent, it causes irritation in 10–20% of users. Tranexamic acid achieves comparable MASI score reduction with a significantly superior safety profile.

2. **Ingredient-to-Product Matching.** Maps indicated ingredients to specific catalogue SKUs, accounting for concentration thresholds and formulation stability. A serum with 0.3% retinol serves a different recommendation context than one with 1.0%, despite both containing the same active ingredient.
3. **Routine Coherence and Temporal Separation.** Constructs a final routine where all products are mutually compatible and appropriately separated into morning (AM) and evening (PM) steps.

8.3 Constraint Categories

The constraint solver enforces four categories of pharmacological safety rules:

- **Hard contraindication constraints:** Clinically dangerous combinations that must never co-occur. Example: L-ascorbic acid paired with copper peptides in the same application step.
- **Temporal separation constraints:** Actives that are individually safe but must be separated into AM/PM routines. Example: Benzoyl peroxide and retinoids.
- **Exfoliation stacking constraints:** Simultaneous use of high-strength AHAs (glycolic acid) and retinoids is flagged due to cumulative barrier stress.
- **Application ordering constraints:** Products must be applied in the correct sequence for optimal absorption. Water-based formulations before oil-based. Active serums before occlusives.

8.4 K-Beauty Ingredient Integration

The SIE's ingredient knowledge graph extends beyond Western dermatological actives to incorporate the "Well-Aging" philosophy of Korean skincare:

- **Galactomyces Ferment Filtrate (GFF).** Activates the aryl hydrocarbon receptor (AHR), upregulating filaggrin expression and strengthening the skin barrier's natural moisturising factor. Recommended when hydration index and TEWL estimation indicate early barrier compromise.
- **Ginsenosides (Panax Ginseng).** Stimulate collagen synthesis through fibroblast activation and inhibit MMP-1, the enzyme primarily responsible for collagen degradation during photoaging. Recommended when micro-wrinkle analysis detects early-stage collagen loss.
- **Snail Mucin (SCA).** Contains glycosaminoglycans and copper peptides that accelerate wound healing and moisture retention. Indicated for barrier recovery scenarios.

9. Capture Quality and the Smartphone Constraint

Deploying clinical analysis on smartphones requires a robust system to handle the extreme variability of real-world imaging conditions. Unlike lab-grade systems like the VISIA Complexion Analysis System, which uses 8x 20MP cameras with cross-polarised flashes, calibrated lighting, and fixed positioning, the SIE must operate under tungsten, fluorescent, and natural light.

9.1 Lighting and Image Quality Assessment (LIQA)

The LIQA module is the first stage of the inference pipeline, evaluating five dimensions of image quality:

1. **Face Detection.** Confirms the presence of a face at sufficient scale, with all key landmarks visible.
2. **Illumination Uniformity.** Measures variance across segmented facial regions to detect shadows that would occlude diagnostically critical features.
3. **Region Resolution.** Evaluates the effective pixel density on the facial surface rather than the raw image resolution.
4. **Sharpness.** Uses Laplacian-based blur detection to identify motion artifacts and out-of-focus conditions.
5. **Exposure Assessment.** Ensures the facial region is within the sensor's dynamic range to preserve textural detail.

Below quality thresholds, the pipeline requests a recapture with **specific guidance** (e.g., "Move to a more evenly lit area"). The impact is significant: **LIQA increases the yield of clinically usable smartphone captures from 32% to 87%**, fundamentally transforming the smartphone from a consumer camera into a calibrated clinical instrument.

9.2 Optical Scattering and Refractive Index Modelling

The quantification of skin texture and hydration levels requires an understanding of how light reflects from a rough surface. The SIE utilises the **Beckmann-Spizzichino reflectance model** to analyse the specular and diffuse components of the reflected light. For a rough surface, the scattered intensity as a function of the incidence and reflectance angles is modelled using the surface slope (m), defined by the surface roughness (σ) and the correlation length (T):

$$I(\theta_i, \theta_r, \phi_r) \approx \frac{2\pi F^2 T^2}{A v_z^2 \sigma^2 \left(1 + \frac{v_{xy}^2 T^2}{v_z^2 \sigma^2}\right)^{3/2}}$$

This model allows the engine to quantify surface roughness (the Texture Index) by analysing the variance in specular reflectance across the facial surface. Furthermore, the hydration index is refined by modelling the refractive index of the stratum corneum, which varies based on water content according to the **law of Gladstone and Dale**:

$$n_{skin} = (1 - c_w) \times 1.5 + c_w \times n_{water}$$

Where **n_water ≈ 1.333**. This allows the SIE to detect the sub-visual gradient of hydration that precedes visible flaking or dryness, enabling preventive intervention before clinical symptoms manifest.

10. Business Intelligence: Regional Demand Sensing

Beyond individual diagnosis, the SIE acts as an **intelligence layer** for the beauty industry by leveraging aggregate, anonymised skin telemetry to drive predictive inventory management and regional demand sensing for enterprise partners.

10.1 Demand Sensing Architecture

The demand-sensing model aggregates skin concern data at the regional, city, and local government area levels. This captures "latent demand": what the population actually needs based on their biology, which is a more powerful signal than purchase history. For example, during the dry Harmattan season in West Africa (November through March), the SIE detects a measurable spike in barrier compromise indicators (elevated TEWL estimation, increased texture index scores, reduced hydration index) across the user population. A 25% increase in barrier compromise detections enables retail partners to proactively shift inventory toward hydrating ceramides, occlusives, and barrier repair formulations.

Label	Category	Scans	Avg Coverage %	Regions	Sample
hyperpigmentation	pigmentation	5095	11.4%	image right cheek, t zone, image left cheek, right cheek, left cheek, forehead, chin	view
comedones	sebum	3993	5.9%	forehead, image right cheek, left cheek, right cheek, image left cheek, chin, t zone	view
enlarged pores	sebum	3301	13.7%	image left cheek, image right cheek, right cheek, left cheek, chin, t zone, forehead	view
under-eye darkness	pigmentation	2523	9.7%	t zone, image left cheek, forehead, image right cheek, chin	view
fine lines	wrinkles	1157	5.9%	forehead, right cheek, image left cheek, image right cheek, chin, left cheek, t zone	view
dullness	hydration	906	26.6%	forehead, t zone, image right cheek, left cheek, chin, image left cheek, right cheek	view
papules	acne	753	7.9%	right cheek, image right cheek, left cheek, chin, image left cheek, forehead, t zone	view
mild shine	sebum	601	18.5%	t zone, forehead, right cheek, left cheek, chin, image left cheek, image right cheek	view
blackheads	sebum	556	4.1%	t zone, left cheek, right cheek, image left cheek, image right cheek, chin	view
under_eye_darkness	pigmentation	460	4.2%	t zone, forehead, image left cheek, image right cheek, chin	view
wrinkles	wrinkles	424	4.3%	chin, forehead, t zone, image left cheek, image right cheek	view
mild dullness	hydration	267	18.3%	forehead, left cheek, chin, right cheek, t zone	view
acne	acne	216	16.3%	left cheek, right cheek, forehead, t zone, image left cheek, chin, image right cheek	view
redness	redness	209	15.7%	image left cheek, right cheek, image right cheek, chin, left cheek, t zone, forehead	view
papule	acne	175	2.1%	image right cheek, image left cheek, left cheek, right cheek, forehead, chin, t zone	view

Fig. 6 — Aggregate Skin Concern Distribution: Population-level analysis across 200,000+ scans showing concern frequency, categorical classification, average coverage percentage, and affected facial regions.

10.2 Commercial Outcomes

Enterprise partners deploying the SIE report the following outcomes:

Metric	Enterprise Outcome
Scan-to-Cart Rate	66%. Two-thirds of users who complete a skin analysis add a recommended product to their cart.
Conversion Lift	1.6x. Users who receive SIE-powered recommendations convert at 1.6 times the rate of standard shoppers.
Return Reduction	30%. Personalised ingredient-condition matching reduces "wrong product" irritation returns.
Usable Capture Yield	87%. LIQA module transforms smartphone capture from 32% usable to 87% usable.

Table 4: Enterprise partner performance metrics across the Oyster partner network.

These commercial outcomes are the direct result of the system's clinical credibility. When a recommendation is based on a 33-dimensional biomarker analysis and 500 pharmacological rules, the user's propensity to purchase and their satisfaction with the product significantly increase.

11. Data Governance and Privacy Architecture

The SIE handles sensitive biometric and health data, necessitating a sophisticated privacy architecture. The system complies with regional data sovereignty requirements, such as Nigeria's NDPR and South Africa's POPIA, through the implementation of **Differential Privacy (DP)**.

11.1 Differential Privacy (DP) Formulation

Differential Privacy ensures that the presence of an individual in the dataset cannot be deduced from the aggregate results. A randomised mechanism K provides (ε, δ)-differential privacy if for all adjacent databases D_1, D_2 and all outputs $S \subseteq$ Range(K):

$$Pr[K(D_1) \in S] \leq exp(\varepsilon) \times Pr[K(D_2) \in S] \; + \; \delta$$

Where **ε** is the privacy budget (controlling the maximum information leakage) and **δ** is the probability of a privacy failure. By applying calibrated Gaussian or Laplace noise to regional trend data, the SIE preserves the statistical validity of population-level health insights while guaranteeing individual anonymity.

11.2 Infrastructure Deployment

The SIE is deployed on a multi-region AWS architecture. The primary African deployment operates from AWS af-south-1 (Cape Town), serving Southern and West African traffic with round-trip latencies significantly lower than routing through European or North American data centres. The multi-tenant architecture isolates partner data at the infrastructure level. On-device preprocessing (LIQA and face detection) runs locally using lightweight mobile-optimised models, ensuring that quality gating operates before any data leaves the device. Only the cropped, compressed facial region at the resolution required by the analysis models is transmitted, reducing payload size and

bandwidth requirements on the bandwidth-constrained mobile networks common in primary markets.

12. Historical Synthesis: From Lustgarten to Systematic AI

In 1910, the Viennese dermatologist **Sigmund Lustgarten**, the first to introduce Salvarsan for syphilis treatment in the United States, advocated for moving dermatology toward an evidence-driven clinical footing. Lustgarten's era relied on "clinical proving": the meticulous manual documentation of visible signs to infer underlying pathology, building a systematic body of knowledge from accumulated observation.

A century later, the Oyster Skin Intelligence Engine represents the computational realisation of Lustgarten's ambition. We have engineered a system that perceives melanin gradients invisible to the human eye, detects barrier compromise from a standard photograph, resolves complex differential diagnoses in under two seconds, and constructs pharmacologically coherent treatment routines from a device that costs less than a stethoscope. Across every skin tone. Skin was never resistant to systematisation; it simply required a system with the computational density to process its optical complexity and the data foundation to represent its global diversity.

13. Conclusion

The Oyster Skin Intelligence Engine marks the convergence of hierarchical vision Transformers, optical physics, and pharmacological rigour. By resolving the socio-technical crisis of representativity through a Melanin-First data strategy and a tone-adaptive inference pipeline conditioned on the Individual Typology Angle, the SIE ensures that advanced dermatological analysis is available to all, regardless of constitutive pigmentation.

The technical innovations detailed in this report include: a Swin Transformer backbone that captures both macro-level skin patterns and micro-level morphological features within a single forward pass; a Swin-Unet segmentation architecture achieving DSC 0.9753; a Kendall uncertainty-weighted multi-task learning framework that maintains precision across 33 simultaneous biomarker quantification tasks; a tone-adaptive inference pipeline conditioned on ITA that ensures clinical accuracy across the full Fitzpatrick spectrum; a deterministic CSP-based recommendation engine enforcing over 500 pharmacological constraints across 18,500+ indexed SKUs; and a LIQA module that transforms the smartphone from a consumer imaging device into a calibrated clinical instrument.

For enterprise partners, these capabilities translate into measurable commercial outcomes: a 66% scan-to-cart rate, a 1.6x conversion lift, and a 30% reduction in product returns. For the beauty industry, the SIE provides the intelligence infrastructure for a transition from demographic-based marketing to condition-based personalisation. For the billions of people whose skin tones were historically treated as edge cases by the technology industry, it provides something more fundamental: a system that sees them with clinical accuracy for the first time.

This is the foundation of a new era of global skin intelligence. Detailed technical reports on individual components of the SIE will be published in the coming months.

References

[1] Groh, M. et al. (2021). "Evaluating Deep Neural Networks Trained on Clinical Images in Dermatology with the Fitzpatrick 17k Dataset." CVPR.

[2] Tschandl, P. et al. (2018). "The HAM10000 Dataset: A Large Collection of Multi-Source Dermatoscopic Images of Common Pigmented Skin Lesions." Scientific Data.

[3] Codella, N. et al. (2018). "Skin Lesion Analysis Toward Melanoma Detection." ISBI Challenge.

[4] Esteva, A. et al. (2017). "Dermatologist-level classification of skin cancer with deep neural networks." Nature.

[5] Liu, Z. et al. (2021). "Swin Transformer: Hierarchical Vision Transformer using Shifted Windows." ICCV.

[6] Brinker, T.J. et al. (2019). "Deep learning outperformed 136 of 157 dermatologists in a head-to-head dermoscopic melanoma image classification task." Annals of Oncology.

[7] Alexis, A.F. et al. (2019). "Acne in Patients with Skin of Color: Practical Management." Journal of Drugs in Dermatology.

[8] Adelekun, A. et al. (2021). "Skin Color in Dermatology Textbooks: An Updated Evaluation and Analysis." Journal of the American Academy of Dermatology.

[9] Marchetti, M.A. et al. (2018). "Results of the 2016 ISIC ISBI challenge." JAMA Dermatology.

[10] Sheu, S.L. (2018). "Treatment of Melasma Using Tranexamic Acid: What's Known and What's Next." American Journal of Clinical Dermatology.

[11] Del Rosario, E. et al. (2018). "Randomized, Placebo-Controlled, Double-Blind Study of Oral Tranexamic Acid in the Treatment of Moderate-to-Severe Melasma." JAAD.

[12] Kendall, A. et al. (2018). "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics." CVPR.

[13] Beckmann, P. & Spizzichino, A. (1987). "The Scattering of Electromagnetic Waves from Rough Surfaces." Artech House.

[14] Internal Oyster Engineering Documentation (2026). "Africa-First Data Strategy and Calibration."

[15] Canfield Scientific. "VISIA Complexion Analysis System." Technical specifications.

[16] Internal Oyster partner analytics (2025). "Retail Conversion and Scan-to-Cart Metrics."

[17] Benedek, T.G. (2004). "The History of Gold Therapy for Tuberculosis." Journal of the History of Medicine.

Works Cited

[1] Bias in Dermatological Datasets: A Critical Analysis of the Underrepresentation of Dark Skin Tones in Melanoma Classification Images. ICCS 2025.

[2] Skin tone distribution in the ISIC dataset, classified by ITA-derived Fitzpatrick type. ResearchGate.

[3] Department of Dermatology History. Icahn School of Medicine at Mount Sinai.

[4] Presentations of Cutaneous Disease in Various Skin Pigmentations: Acne Vulgaris. PMC.

[5] Disparities in dermatology educational resources. ResearchGate.

[6] Visia 3D. Canfield Scientific.

[7] How Foxtale used AI Skin Analysis to boost conversions and engagement with GlamAR.

[8] Liu, Z. et al. (2021). "Swin Transformer: Hierarchical Vision Transformer using Shifted Windows." ICCV.