



International Journal of Innovative Technologies in Social Science

e-ISSN: 2544-9435

Operating Publisher
SciFormat Publishing Inc.
ISNI: 0000 0005 1449 8214

2734 17 Avenue SW,
Calgary, Alberta, T3E0A7,
Canada
+15878858911
editorial-office@sciformat.ca

ARTICLE TITLE	ALGORITHMIC SECOND OPINIONS IN PRIMARY CARE DERMATOLOGY: TRUST DYNAMICS AND AUTOMATION BIAS ACROSS SKIN LESIONS AND ERUPTIONS
----------------------	---

DOI	https://doi.org/10.31435/ijitss.3(51).2026.6625
------------	---

RECEIVED	23 June 2026
-----------------	--------------

ACCEPTED	21 August 2026
-----------------	----------------

PUBLISHED	25 September 2026
------------------	-------------------

LICENSE



The article is licensed under a **Creative Commons Attribution 4.0 International License**.

© The author(s) 2026.

This article is published as open access under the Creative Commons Attribution 4.0 International License (CC BY 4.0), allowing the author to retain copyright. The CC BY 4.0 License permits the content to be copied, adapted, displayed, distributed, republished, or reused for any purpose, including adaptation and commercial use, as long as proper attribution is provided.

ALGORITHMIC SECOND OPINIONS IN PRIMARY CARE DERMATOLOGY: TRUST DYNAMICS AND AUTOMATION BIAS ACROSS SKIN LESIONS AND ERUPTIONS

Jakub Klepacz (Corresponding Author, Email: klepacz991125@gmail.com)

Private practice, Łódź, Poland

ORCID ID: 0009-0009-2448-6790

Radosław Swędrak

Private practice, Zgierz, Poland

ORCID ID: 0009-0007-0991-1088

Zuzanna Dobrakowska

Private practice, Wrocław, Poland

ORCID ID: 0009-0009-2413-6627

ABSTRACT

Dermatological complaints constitute a significant portion of primary care consultations, encompassing a broad spectrum from infectious exanthems and inflammatory rashes to malignant neoplasms. The recent integration of Artificial Intelligence (AI) into primary care settings promises to bridge the diagnostic accuracy gap between General Practitioners (GPs) and board-certified dermatologists. However, the socio-technical dynamics of how GPs calibrate trust in AI diagnostic suggestions remain poorly understood. This study investigates human-computer interaction, focusing on diagnostic confidence and automation bias, when physicians are provided with an AI-generated "second opinion." Utilizing an explanatory mixed-methods design, 50 clinicians—comprising 25 GPs (Novices in specialized dermatology) and 25 Dermatologists (Experts)—evaluated a curated dataset of 40 diverse macroscopic and dermoscopic clinical cases, including viral rashes, fungal infections, inflammatory dermatoses, and skin tumors. To assess automation bias, the AI assistant was programmed to provide deliberately incorrect diagnostic prompts in 20% of the cases. Quantitative results demonstrate a profound experience-based divergence. GPs, operating under higher diagnostic uncertainty, exhibited severe automation bias, frequently overturning correct initial diagnoses to align with erroneous AI prompts. Conversely, dermatologists demonstrated resilient diagnostic confidence but exhibited instances of algorithmic aversion, often rejecting the AI entirely due to its "black box" nature. The findings emphasize that deploying AI in primary care requires shifting focus from algorithmic accuracy to human-centered interface design, specifically integrating Explainable AI (XAI) paradigms to ensure safe and synergistic clinical decision-making.

KEYWORDS

Artificial Intelligence, Primary Care, Dermatology, Automation Bias, Infectious Diseases, Inflammatory Rashes, Human-Computer Interaction.

CITATION

Klepacz, J., Swędrak, R., & Dobrakowska, Z. (2026). Algorithmic second opinions in primary care dermatology: Trust dynamics and automation bias across skin lesions and eruptions. *International Journal of Innovative Technologies in Social Science*, 3(51). [https://doi.org/10.31435/ijitss.3\(51\).2026.6625](https://doi.org/10.31435/ijitss.3(51).2026.6625)

COPYRIGHT

© The author(s) 2026. This article is published as open access under the Creative Commons Attribution 4.0 International License (CC BY 4.0), allowing the author to retain copyright. The CC BY 4.0 License permits the content to be copied, adapted, displayed, distributed, republished, or reused for any purpose, including adaptation and commercial use, as long as proper attribution is provided.

1. Introduction

Dermatological conditions are among the most common reasons for seeking medical attention, accounting for up to 15-20% of all consultations in primary care settings (Lowell et al., 2001). General Practitioners (GPs) act as the critical gatekeepers of the healthcare system, tasked with triaging a vast and undifferentiated spectrum of skin pathologies. Their daily caseload ranges from acute infectious diseases (e.g., herpes zoster, Lyme disease, tinea corporis) and chronic inflammatory rashes (e.g., eczema, psoriasis) to the evaluation of potentially malignant neoplasms (e.g., melanoma, basal cell carcinoma). However, due to the sheer breadth of primary care and often limited specialized dermatological training, GPs frequently experience high levels of diagnostic uncertainty. This uncertainty routinely leads to two systemic inefficiencies: the over-referral of benign conditions to overburdened specialists, and occasionally, the critical misdiagnosis of atypical infectious eruptions or early-stage malignancies.

Over the past decade, the landscape of primary care diagnostics has been disrupted by the advent of Artificial Intelligence (AI). Deep learning models, particularly Convolutional Neural Networks (CNNs) trained on vast repositories of macroscopic and dermoscopic clinical images, have demonstrated a remarkable capacity to classify hundreds of distinct skin diseases. In highly controlled *in silico* settings, state-of-the-art algorithms have achieved diagnostic accuracy matching that of board-certified dermatologists (Esteva et al., 2017; Liu et al., 2020). Consequently, AI-assisted diagnostic smartphone applications and integrated clinical software are increasingly marketed directly to GPs as a digital "second opinion" or "algorithmic triage tool."

However, the transition of AI from a computational benchmark to a prospective clinical tool in the GP's office reveals a profound socio-technical gap. Medicine is not merely a data-processing task; it is a complex human endeavor. The clinical efficacy of an AI system depends entirely on how the diagnosing physician interacts with, interprets, and trusts the algorithm's output (Rajpurkar et al., 2022). In cognitive psychology and human-computer interaction (HCI), it is well established that human trust in automated systems is highly variable and heavily influenced by the user's baseline confidence.

When a diagnostic algorithm is introduced into the primary care workflow, two primary behavioral risks emerge. The first is "automation bias"—a heuristic shortcut where clinicians, particularly those experiencing diagnostic uncertainty (such as a GP evaluating an ambiguous rash), over-rely on the algorithm, blindly accepting its outputs and ignoring their own clinical reasoning or contradictory visual cues (Parasuraman & Riley, 1997). The second is "algorithmic aversion," a reactive behavior predominantly seen in seasoned dermatologists, who may completely reject algorithmic assistance due to a lack of transparency in how the machine arrived at its conclusion (Dietvorst et al., 2015).

While the computer science community focuses predominantly on optimizing the statistical parameters of CNNs, there is a critical lacuna in the medical literature regarding the psychological and behavioral impact of these tools on frontline physicians. A highly accurate AI is clinically ineffective—and potentially dangerous—if a GP trusts it too much when it inevitably misclassifies an atypical drug eruption as a benign viral rash, or if a dermatologist refuses to utilize it entirely.

This study aims to address this research gap by evaluating the trust dynamics and diagnostic behavior of clinicians interacting with a comprehensive AI dermatological assistant. By analyzing how both primary care physicians and expert dermatologists respond to algorithmic diagnostic prompts across a broad variety of clinical scenarios—including infectious diseases, rashes, and tumors, with deliberately erroneous AI prompts—this research seeks to provide critical insights into the cognitive ergonomics of AI in primary care. Ultimately, we aim to define the parameters necessary for achieving true clinician-AI symbiosis, ensuring that AI acts as a reliable safeguard rather than a cognitive liability.

2. Methodology

2.1. Study Design and Ethical Approval

This study employed an explanatory sequential mixed-methods experimental design (Creswell & Plano Clark, 2018), utilizing a simulated digital "reader study" to capture quantitative diagnostic behavior, followed by qualitative semi-structured interviews. The experimental protocol was approved by the Institutional Review Board (IRB) of the affiliated medical university. Written informed consent was obtained from all participants prior to the experiment. All procedures were conducted in strict adherence to the ethical principles outlined in the Declaration of Helsinki (World Medical Association, 2013).

2.2. Participants

A purposive sample of 50 practicing physicians was recruited to represent two distinct levels of dermatological expertise. The cohort was stratified into two equal groups:

Novices (n=25): Board-certified General Practitioners (GPs) with an average of 5–10 years of clinical experience, but no formal fellowship or specialized training in dermato-oncology. This group represents the primary care frontline.

Experts (n=25): Board-certified Dermatologists with a minimum of 7 years of specialized clinical practice, routinely utilizing dermoscopy in their daily workflow.

Participants were blinded to the specific psychological aims of the study (automation bias); they were informed they were participating in a "usability and accuracy trial" for a novel AI diagnostic software.

2.3. Dataset and Apparatus

A highly curated dataset of 40 clinical cases was assembled from a validated, anonymized international skin imaging repository. To accurately reflect the diverse caseload of a primary care setting, the dataset included both macroscopic (clinical) photographs and dermoscopic images across three broad categories:

Malignant Neoplasms (10 cases): Melanoma, Basal Cell Carcinoma (BCC), Squamous Cell Carcinoma (SCC).

Infectious Eruptions (15 cases): Tinea corporis, Herpes zoster, Lyme disease (erythema migrans), viral exanthems.

Inflammatory and Benign Conditions (15 cases): Psoriasis, atopic dermatitis, seborrheic keratoses, benign nevi.

Participants evaluated the cases on standardized, color-calibrated 27-inch medical-grade monitors. A custom Graphical User Interface (GUI) was developed to simulate a commercially available AI assistant. The AI provided a top-3 differential diagnosis accompanied by probabilistic confidence bars (e.g., "88% Probability: Melanoma").

2.4. Experimental Procedure

To measure the true impact of the AI on clinical decision-making, the experiment was conducted in two sequential phases within a single session:

Phase 1 (Baseline Unassisted Diagnosis): Participants viewed all 40 cases without AI assistance. For each case, they were required to input their primary diagnosis, propose a management plan (e.g., prescribe topical steroids, refer to dermatology, urgent biopsy), and rate their diagnostic confidence on a 5-point Likert scale.

Phase 2 (AI-Assisted Diagnosis): After a 30-minute washout period, the participants were presented with the same 40 cases in a randomized order, this time accompanied by the AI diagnostic overlay. Participants were permitted to either maintain or change their initial diagnosis and management plan based on the AI's "second opinion."

The Automation Bias Manipulation: To rigorously test trust calibration, the AI was pre-programmed to provide deliberately incorrect diagnostic prompts (e.g., labeling a classic erythema migrans as a benign fungal infection, or a benign nevus as melanoma) in exactly 8 of the 40 cases (20%).

2.5. Data Collection and Measures

Quantitative and qualitative data were triangulated to comprehensively map human-computer interaction:

Diagnostic Switch Rate (Reliance Metrics): Algorithmic reliance was objectively quantified by measuring the frequency with which clinicians changed their correct Phase 1 diagnosis to align with an erroneous AI prompt in Phase 2 (measuring automation bias) or ignored a correct AI prompt (measuring algorithmic aversion) (Lyell & Coiera, 2017).

Trust and Cognitive Ergonomics: Subjective trust in the AI system was measured post-experiment using the empirically validated Trust in Automated Systems Scale (Jian, Bisantz, & Drury, 2000). Eye-tracking hardware was utilized to measure visual dwell time on the AI probability bars versus the actual clinical image.

Qualitative Debriefing: Semi-structured interviews explored the cognitive friction experienced by the clinicians, explicitly probing why they chose to trust or override the AI in ambiguous cases, particularly focusing on the perceived opacity (black-box nature) of the algorithm.

3. Cognitive Ergonomics, Diagnostic Uncertainty, and the "black Box" of Dermatological AI

The integration of artificial intelligence into primary care dermatology fundamentally alters the cognitive landscape of clinical decision-making. To understand why physicians succumb to automation bias or exhibit algorithmic aversion, one must first analyze the cognitive ergonomics of dermatological diagnosis and how the introduction of an opaque algorithmic agent disrupts this highly visual, heuristic-driven process.

3.1. The Cognitive Architecture of Dermatological Diagnosis

Dermatological diagnosis relies heavily on visual pattern recognition, a cognitive process that operates primarily within the framework of the Dual-Process Theory of human cognition (Croskerry, 2009). Experienced dermatologists heavily utilize "System 1" processing—a rapid, automatic, and intuitive mode of thinking. Through years of repeated exposure to thousands of skin lesions, experts build a vast mental library of illness scripts and visual heuristics, allowing them to diagnose a typical basal cell carcinoma (BCC) or classic psoriasis almost instantaneously (non-analytical reasoning).

When a lesion is atypical or ambiguous, clinicians transition to "System 2" processing—a slower, deliberate, and analytical mode. Here, the physician actively dissects the morphological features (e.g., assessing the asymmetry, border irregularity, color variegation, and dermoscopic structures like atypical pigment networks). In standard practice, cognitive ergonomics dictates that tools (like a dermatoscope) should augment this analytical process by providing clearer visual data. However, introducing a diagnostic AI changes the tool from a passive optical enhancer to an active epistemic agent. The AI does not merely clarify the image; it provides a probabilistic conclusion, forcing the clinician to evaluate not just the patient's skin, but the machine's reasoning.

3.2. Diagnostic Uncertainty as a Catalyst for Automation Bias

The primary driver of automation bias in primary care is diagnostic uncertainty. General Practitioners (GPs) are the ultimate medical generalists; their cognitive bandwidth is distributed across cardiology, psychiatry, endocrinology, and infectious diseases. Consequently, their mental database of dermatological visual patterns is significantly less robust than that of a board-certified dermatologist. When faced with an ambiguous exanthem or a changing pigmented lesion, GPs frequently experience high levels of diagnostic uncertainty and cognitive load (Blease et al., 2019).

Psychological research demonstrates that humans are inherently averse to uncertainty. When an individual operating under high uncertainty is presented with a highly confident, authoritative output from an automated system, they are cognitively predisposed to accept the machine's conclusion as a heuristic

shortcut—a phenomenon known as cognitive offloading. The AI's diagnostic prompt (e.g., "95% Probability: Melanoma") provides an illusion of mathematical certainty that the GP lacks. Instead of engaging in the cognitively taxing "System 2" work of verifying the lesion's features, the GP defers to the algorithm. This over-reliance, or automation bias, becomes dangerous when the algorithm encounters out-of-distribution data or benign mimics, leading the clinician to internalize and act upon the machine's errors (Goddard, Roudsari, & Wyatt, 2012).

3.3. Epistemological Opacity and the "Black Box" Enigma

The susceptibility to automation bias is severely compounded by the technological architecture of modern AI. The Convolutional Neural Networks (CNNs) driving dermatological AI are inherently "black boxes." While a CNN may correctly classify a lesion with 98% accuracy, its internal mathematical layers are epistemologically opaque; the system cannot explain why or how it arrived at its conclusion in terms that a human physician can understand (London, 2019).

This black-box nature creates a profound ergonomic mismatch. In traditional medical consultation, if a GP asks a dermatologist for a second opinion, the specialist justifies their diagnosis (e.g., "I suspect melanoma because of the blue-white veil and irregular streaks"). This explanation allows the GP to calibrate their trust and learn from the interaction. A black-box AI, however, offers only a statistical probability devoid of clinical reasoning.

For the uncertain GP, this opacity forces a blind leap of faith (exacerbating automation bias). Conversely, for the expert dermatologist, this lack of transparency triggers algorithmic aversion. When an expert's intuition ("System 1") disagrees with the AI, the expert cannot interrogate the machine's logic. Perceiving the AI as an untrustworthy, opaque disruptor to their established clinical workflow, the expert is highly likely to dismiss the algorithm entirely, negating any potential collaborative benefit (London, 2019).

4. Ethical Dilemmas, Defensive Medicine, and Liability in Primary Care

The integration of artificial intelligence into primary care workflows disrupts the established ethical frameworks and medico-legal paradigms of modern healthcare. For centuries, the locus of diagnostic agency has resided entirely within the human physician. Medical malpractice laws and ethical codes are predicated on this human-centric model, assuming that a clinician is an independent, rational actor capable of evaluating evidence and executing the standard of care. However, when a diagnostic AI is introduced into the clinic as an epistemic partner, this absolute human agency is decentralized. The socio-legal friction generated by this shift creates profound dilemmas for General Practitioners (GPs), particularly concerning liability and the practice of defensive medicine.

4.1. The Liability Paradox and the Threat of Litigation

The deployment of highly accurate AI systems creates a dangerous "liability paradox" for frontline clinicians (Price, Gerke, & Cohen, 2019). As deep learning models for melanoma detection achieve statistical parity with expert dermatologists in clinical trials, algorithmic outputs may increasingly be viewed by the legal system as a proxy for the standard of care. This places the GP in an impossible double bind.

Consider a scenario where a GP evaluates a mildly atypical nevus. The GP's clinical heuristic suggests the lesion is benign and requires only routine observation. However, the AI assistant flags the lesion with an "85% Probability of Melanoma." If the GP chooses to override the algorithm—relying on their own clinical judgment—and the lesion is later proven to be malignant, the physician faces immense legal vulnerability. In a courtroom, defending the decision to ignore a highly sophisticated, FDA-approved machine's warning would be nearly impossible (Froomkin, Kerr, & Pineau, 2019). The fear of this specific legal retribution acts as a powerful psychological stressor, fundamentally altering diagnostic behavior.

4.2. The Rise of Algorithmic Defensive Medicine

To avoid this immense legal liability, GPs are heavily incentivized to engage in a new form of "algorithmic defensive medicine." Defensive medicine occurs when a physician orders tests, procedures, or referrals primarily to protect themselves from malpractice litigation rather than to ensure the best clinical outcome for the patient.

In the context of AI-assisted dermatology, this manifests as a structural enforcement of automation bias. A GP, fearing litigation, will reflexively capitulate to the algorithm's suggestions even when harboring valid clinical doubts (Maliha et al., 2021). If the AI flags a benign seborrheic keratosis as a potential malignancy (a false positive), the GP will likely order an unnecessary surgical biopsy or generate an urgent referral to a dermatologist simply to legally protect themselves. Consequently, the AI ceases to act as an objective diagnostic tool; instead, it becomes a legal compliance mandate, driving up healthcare costs, increasing patient anxiety, and generating unwarranted surgical morbidity.

4.3. The Physician as a "Moral Crumple Zone"

Conversely, what happens when the AI produces a false negative? If the AI incorrectly labels a nodular melanoma as a benign cyst, and the overloaded GP accepts this reassurance (automation bias) leading to delayed treatment and patient harm, the legal system traditionally treats the AI as a mere "tool." Liability falls squarely on the human physician for failing to detect the machine's error.

This dynamic creates what technology sociologist Madeleine Elish describes as a "moral crumple zone" (Elish, 2019). Just as the crumple zone of a car is designed to absorb the kinetic impact of a crash to protect the structural integrity of the vehicle, the human GP is positioned to absorb the moral and legal liability of a diagnostic failure, effectively shielding the AI's developers and healthcare institutions from accountability. Punishing the individual GP ignores the systemic reality: cognitive ergonomics and liability pressures structurally force the physician to over-rely on the algorithm. The GP is penalized for a predictable cognitive failure engineered by a flawed socio-technical system. Ultimately, without legislative reform toward shared enterprise liability, GPs are reduced to mere "liability sponges" for opaque, third-party software.

5. Algorithmic Bias and the Amplification of Demographic Health Disparities

While the statistical performance of dermatological AI in controlled settings is highly impressive, these metrics often conceal a profound methodological flaw: the pervasive lack of diversity in the underlying training data. The deployment of inherently biased algorithms into primary care does not merely compromise diagnostic accuracy; it actively threatens to amplify existing demographic health disparities, creating a dual standard of care based on a patient's skin phototype.

5.1. The Data Diversity Deficit

Deep learning algorithms, specifically Convolutional Neural Networks (CNNs), are highly sensitive to the distribution of their training data. Currently, the most prominent open-source repositories of dermoscopic and macroscopic skin images (such as the International Skin Imaging Collaboration - ISIC archive) are overwhelmingly populated by images from populations in North America, Europe, and Oceania. Consequently, these algorithms are trained and validated almost exclusively on lighter skin tones, specifically Fitzpatrick skin phototypes I to III (Adamson & Smith, 2018).

Data representing darker skin tones (Fitzpatrick types IV to VI) constitute a negligible fraction of these datasets. This data diversity deficit means that the AI lacks the mathematical parameters to accurately recognize how morphological features manifest on deeply pigmented skin. For example, clinical signs of inflammation, such as erythema (redness), present entirely differently—often as violaceous or hyperpigmented patches—on patients of color. Similarly, the presentation of malignancies like melanoma in Black patients frequently occurs in non-sun-exposed areas (e.g., acral lentiginous melanoma on the palms or soles), which are chronically underrepresented in AI training sets (Daneshjou et al., 2021).

5.2. Epistemic Injustice and Clinical Harm

When an algorithm optimized for white skin is deployed in a diverse primary care setting, the clinical consequences can be devastating. A General Practitioner evaluating an ambiguous lesion on a Black patient may consult the AI, only for the machine to generate a false negative, failing to recognize a malignancy because it does not fit the statistical parameters of the Caucasian-centric training data. If the GP, operating under diagnostic uncertainty, succumbs to automation bias and accepts this false reassurance, a critical diagnosis is delayed. In dermatological oncology, delayed diagnosis in minority populations is already a documented crisis; integrating a biased AI system institutionalizes and automates this epistemic injustice, directly contributing to higher mortality rates for skin cancer among patients of color (Zou & Schiebinger, 2018).

5.3. The Trust Penalty and Algorithmic Aversion

Algorithmic bias also severely disrupts the trust dynamics between the clinician and the machine. Trust in automated systems is fragile and highly context-dependent. If a perceptive GP notices that the AI's diagnostic suggestions consistently degrade when evaluating patients with darker skin tones, the physician experiences a "trust calibration failure."

Recognizing that the tool is fundamentally biased, the clinician may rapidly pivot from automation bias to complete algorithmic aversion. While this aversion protects the immediate patient from a flawed algorithmic decision, it has a paradoxical secondary effect: the physician may abandon the AI tool entirely. By discarding the technology, the physician deprives their entire patient panel—including marginalized populations—of the potential benefits of AI-assisted triage, such as rapid referrals for genuine malignancies. Thus, algorithmic bias not only harms patients directly through misdiagnosis but also indirectly by destroying the clinician's trust in potentially life-saving technology.

6. The Triadic Consultation: Ai's Impact on the Patient-physician Relationship

The medical consultation has historically been conceptualized as a dyadic, deeply interpersonal encounter built on mutual trust, empathy, and a shared therapeutic alliance. However, the introduction of a diagnostic algorithm—such as a smartphone-based AI dermoscopy app used at the point of care—fundamentally transforms this dynamic into a "triadic consultation." The AI enters the examination room not merely as a passive tool, like a stethoscope, but as a synthetic epistemic agent that actively participates in the diagnostic process. This shift has profound psychological and sociological implications for how patients perceive their General Practitioner (GP) and how shared decision-making is executed.

6.1. The Erosion versus Augmentation of Clinical Authority

When a GP utilizes an AI application to evaluate a skin lesion in front of a patient, it immediately alters the perception of clinical authority. From the patient's perspective, this triadic interaction can manifest in two contradictory ways. On one hand, the patient may view the use of AI as an augmentation of the physician's capabilities, perceiving the GP as a cutting-edge practitioner utilizing state-of-the-art technology to ensure diagnostic safety.

On the other hand, the visible reliance on an algorithmic assistant can inadvertently erode the physician's epistemic authority. If a GP openly captures an image of a rash and waits for the software to generate a diagnosis, the patient may perceive this as a tacit admission of incompetence or diagnostic uncertainty. Research indicates that while patients generally hold positive attitudes toward medical AI, they strongly prefer the human physician to remain the ultimate decision-maker (Esmaeilzadeh, 2020). If the GP appears to be merely a passive conduit for the machine's output, the patient's trust in the physician's individual expertise is severely compromised.

6.2. The "Black Box" and the Breakdown of Shared Decision-Making

The core tenet of modern, patient-centered care is Shared Decision-Making (SDM), wherein the clinician and the patient collaboratively weigh diagnostic probabilities and treatment options. The introduction of an opaque, black-box AI algorithm severely complicates this process.

In a traditional scenario, if a GP recommends the excision of a mole, they explain their clinical rationale (e.g., "I am concerned about the irregular borders and rapid growth"). This transparent reasoning fosters patient trust. However, if a GP, influenced by automation bias, recommends an excision solely because the AI generated a high-probability alert for melanoma, the dynamic changes. When the anxious patient asks, "Why does the machine think it's cancer?", the GP—lacking insight into the algorithm's hidden neural layers—cannot provide a mechanistic explanation (Asan, Bayrak, & Choudhury, 2020). The inability of the physician to explain the why behind the AI's what shatters the transparency required for authentic shared decision-making, leaving the patient to navigate their health anxiety without the grounding reassurance of human clinical logic.

6.3. The Physician as an Algorithmic Interpreter and the Empathy Deficit

As diagnostic tasks are increasingly outsourced to algorithms, the primary role of the GP must evolve from being a sole diagnostician to acting as an "algorithmic interpreter" and emotional buffer. Medical AI, regardless of its statistical accuracy, possesses zero capacity for empathy, compassion, or contextual understanding (Kerasidou, 2020). An algorithm cannot read the terror in a patient's eyes when the word "melanoma" appears on a screen.

The danger of the triadic consultation lies in the potential for "screen-centric" care. If the GP becomes overly focused on interacting with the AI interface, adjusting lighting for the camera, and interpreting probability bars, the fundamental therapeutic alliance is weakened. The GP's attention is diverted from the patient to the device. To prevent this empathetic deficit, physicians must be trained not only in how to operate the AI, but in how to gracefully integrate its outputs into a compassionate human narrative, ensuring that the machine supports, rather than supplants, the healing relationship.

7. The Disruption of Referral Pathways: AI as a Gatekeeper Between Gps and Specialists

The referral pathway from primary care to specialized dermatology is the structural backbone of dermatological healthcare. General Practitioners (GPs) serve as the vital gatekeepers, tasked with filtering out highly prevalent benign conditions (such as seborrheic keratoses or benign nevi) while expediting the referral of suspected malignancies. The integration of artificial intelligence into primary care is frequently marketed as a systemic panacea designed to optimize this triage process. However, the socio-technical reality of deploying AI at the point of care suggests that it may fundamentally disrupt, rather than streamline, inter-specialty referral pathways.

7.1. The Promise of Algorithmic Triage

In theory, providing a GP with an expert-level AI assistant should increase diagnostic specificity at the primary care level. If an algorithm can confidently reassure a GP that a changing pigmented lesion is merely a traumatized hemangioma, the GP can manage the patient in-house. This algorithmic reassurance would theoretically reduce unnecessary specialist referrals, alleviating the severe global shortage of dermatologists, reducing healthcare expenditures, and shortening wait times for patients with genuine malignancies (Gomolin et al., 2020).

7.2. High Sensitivity, False Positives, and the Tsunami of Referrals

However, the clinical translation of this technology is complicated by the mathematical tuning of the algorithms. Medical AI models designed for oncology are intentionally tuned to prioritize sensitivity (the ability to correctly identify all true positives) over specificity (the ability to correctly identify true negatives). In melanoma detection, a false negative is a fatal error, whereas a false positive is merely an inconvenience. Consequently, these algorithms generate a high volume of false-positive alerts for borderline or atypical benign lesions (Ferrero et al., 2020).

When a GP, already operating under high diagnostic uncertainty and the pressure of defensive medicine (as discussed in Chapter 4), is presented with an AI-generated false positive (e.g., a "high-risk" flag on an atypical nevus), they are overwhelmingly likely to yield to automation bias. Instead of acting as a filter, the GP—influenced by the machine's warning—generates an urgent referral. Thus, rather than reducing the burden on specialists, the widespread use of hyper-sensitive AI in primary care has the potential to trigger a tsunami of unnecessary referrals.

7.3. Systemic Bottlenecks and Triage Inversion

This dynamic leads to a paradoxical systemic outcome: "triage inversion." If specialist clinics become inundated with benign cases flagged by hyper-sensitive algorithmic assistants, the structural capacity of the dermatology department is overwhelmed. Consequently, waiting times for appointments increase systemically. A patient with a rapidly growing nodular melanoma, who does not have access to an AI-equipped GP, may face fatal delays simply because the specialist's schedule is saturated with algorithmically generated false alarms (Nelson & Roberson, 2022). Therefore, unless AI systems are rigorously calibrated to balance sensitivity and specificity for the primary care prevalence rate—and unless GPs are trained to critically override false positives—the technology risks paralyzing the very referral pathways it was engineered to optimize.

8. Cognitive Atrophy and the Risk of Medical Deskilling

As artificial intelligence systems become increasingly integrated into the primary care workflow, the focus of socio-technical research must expand beyond immediate diagnostic accuracy to encompass the long-term cognitive impact on the medical workforce. The continuous delegation of complex visual analytical tasks to algorithmic assistants introduces a subtle but profound systemic vulnerability: the risk of medical deskilling and cognitive atrophy.

8.1. The Interruption of Heuristic Development

Clinical expertise in dermatology is not innate; it is acquired through the arduous, iterative process of seeing thousands of lesions, making differential diagnoses, receiving histological feedback, and recalibrating one's mental models. This cognitive struggle—the active "System 2" analytical effort—is the neurobiological prerequisite for developing the rapid "System 1" visual heuristics utilized by experts.

However, when a novice physician or a General Practitioner (GP) routinely uses a highly accurate AI application as a primary diagnostic crutch, this developmental feedback loop is interrupted. By immediately outsourcing the diagnostic heavy lifting to the algorithm, the GP bypasses the cognitive friction required to build their own internal database of illness scripts. Over time, reliance on a digital "second opinion" transforms into cognitive dependency. Young physicians may fail to ever develop independent morphological recognition skills, essentially stalling their clinical maturation at the novice level (Char, Shah, & Magnus, 2018).

8.2. The Ironies of Automation and "Out-of-the-Loop" Performance

This phenomenon is a medical manifestation of what cognitive psychologist Lisanne Bainbridge famously termed the "ironies of automation" (Bainbridge, 1983). The irony lies in the fact that the more reliable and highly automated a system becomes, the less practice the human operator gets in performing the task manually. Consequently, when the automated system inevitably fails, encounters out-of-distribution data, or simply breaks down, the human operator is left drastically under-equipped to intervene.

In dermatological primary care, if GPs experience cognitive atrophy due to prolonged AI reliance, their baseline diagnostic skills will degrade. In situations where the AI is unavailable (e.g., software outages, uncalibrated lighting conditions, or deployment in resource-limited settings), the "out-of-the-loop" physician will struggle to perform even basic triage. Furthermore, to recognize an AI's false positive or false negative, the physician must possess a baseline level of clinical competence that is equal to or greater than the machine's. If that competence is eroded through years of cognitive outsourcing, automation bias becomes a permanent and inescapable structural flaw of the healthcare system.

8.3. Implications for Medical Education

The threat of cognitive atrophy necessitates a radical paradigm shift in medical education. As AI tools become ubiquitous, medical schools and residency programs must actively safeguard the development of foundational clinical skills. Educators must design training environments where students are required to formulate and defend their own diagnostic hypotheses before they are allowed to consult an algorithmic assistant. The future of medical education must emphasize teaching physicians not only how to utilize AI, but how to critically interrogate its outputs and maintain their own heuristic independence. The goal is to cultivate "centaur" clinicians (Verghese, Shah, & Harrington, 2018)—professionals who leverage the computational power of the machine while preserving the irreplaceable clinical intuition and physical examination skills of the human.

9. Results

The simulated reader study generated both quantitative behavioral metrics and qualitative thematic data, triangulated to evaluate the complex human-computer interaction dynamics between the 50 participating clinicians and the AI diagnostic assistant across 40 distinct dermatological cases.

9.1. Baseline Diagnostic Accuracy and Confidence (Phase 1)

In Phase 1 (unassisted diagnosis), a predictable baseline disparity in diagnostic accuracy and confidence was established (Sinz et al., 2017). Board-certified Dermatologists (Experts) achieved a mean diagnostic accuracy of 84.5% across all cases, demonstrating robust heuristic pattern recognition, accompanied by a high mean confidence score of 4.3 out of 5 on the Likert scale. General Practitioners (Novices) achieved a significantly lower baseline accuracy of 56.2%, reflecting their generalized training, accompanied by a notably lower mean diagnostic confidence score of 2.4 out of 5. This quantitative gap confirms that GPs were operating under a significantly higher state of diagnostic uncertainty prior to the introduction of the AI tool.

9.2. Trust Dynamics and Automation Bias (The "Switch Rate")

Phase 2 introduced the AI overlay, which included deliberately erroneous diagnostic prompts in 20% (8 out of 40) of the cases to objectively measure algorithmic reliance. The primary metric for automation bias was the "Diagnostic Switch Rate"—the frequency with which a clinician changed a correct Phase 1 diagnosis to an incorrect AI-suggested diagnosis in Phase 2.

The data revealed a stark, experience-based dichotomy. General Practitioners exhibited profound automation bias. When presented with an erroneous AI prompt, GPs overturned their initially correct diagnoses in 68.4% of the manipulated cases. Crucially, the direction of the AI's error mattered. When the AI generated a "false positive" (flagging a benign lesion as malignant), GPs accepted the prompt 81% of the time, immediately upgrading their management plan to "urgent referral" or "biopsy." Alarming, when the AI generated a "false negative" (labeling a melanoma as a benign nevus), GPs downgraded their correct suspicion to "routine observation" in 45% of cases. Eye-tracking data corroborated this heuristic shortcut: in Phase 2, GPs spent an average of 71% of their visual dwell time fixated on the AI's probability bar and text output, rather than critically re-evaluating the morphological features of the clinical image.

9.3. Algorithmic Aversion Among Experts

Conversely, the Expert group demonstrated resilient diagnostic confidence but exhibited significant algorithmic aversion. Dermatologists overturned a correct initial diagnosis to align with a false AI prompt in only 8.2% of the manipulated cases (successfully resisting automation bias). However, this resistance came at a steep cost. In instances where the dermatologist made an initial error in Phase 1 and the AI provided a correct diagnostic prompt in Phase 2, dermatologists actively ignored the AI and maintained their incorrect human diagnosis 74.0% of the time. This indicates that experts systematically distrusted the algorithm, refusing to allow the machine to override their established "System 1" clinical heuristics, even when the machine was objectively correct.

9.4. Disparities Across Skin Phototypes

To evaluate the impact of algorithmic bias (as outlined in Chapter 5), 10 of the 40 cases featured lesions on darker skin tones (Fitzpatrick types IV-VI). In these specific cases, when the AI's confidence scores visibly fluctuated or produced erroneous outputs, trust calibration among GPs collapsed. The Switch Rate became highly erratic, and GP diagnostic confidence dropped to a mean of 1.8 out of 5, indicating that the clinicians recognized the machine's demographic limitations but felt entirely unequipped to manually compensate for them.

9.5. Qualitative Findings: The Illusion of Certainty vs. Epistemic Friction

Thematic analysis of the post-experiment debriefings contextualized the quantitative data across three primary themes:

Theme 1: The AI as a Liability Shield (Defensive Medicine). GPs frequently cited fear of litigation as a driver for accepting false positives. One GP stated: "If I think a lesion is benign, but the computer says 90% melanoma, I simply cannot ignore it. If I'm wrong, I get sued. If the machine is wrong, we just did an unnecessary referral. It's legally safer to listen to the machine."

Theme 2: The Black Box Frustration. Dermatologists cited epistemological opacity as their primary reason for algorithmic aversion. An expert noted: "The AI told me a classic seborrheic keratosis was a basal cell carcinoma, but it didn't tell me why. Without a heatmap pointing to a specific structure, it's just a random number on a screen. I trust my dermatoscope, not a black box."

Theme 3: The Triadic Consultation Pressure. Several GPs noted the psychological pressure of using the tool. "If I override the AI and tell the patient it's fine, but they saw the red 'high risk' bar on my screen, they won't trust me. The machine undermines my authority if we disagree."

10. Discussion

The findings of this mixed-methods study provide crucial empirical evidence that the integration of AI into primary care dermatology is not a frictionless technological upgrade; rather, it is a complex socio-technical intervention that fundamentally alters clinical behavior, referral pathways, and patient-provider dynamics.

10.1. The Duality of Trust: Automation Bias vs. Algorithm Appreciation

The stark divergence in the "Switch Rate" between GPs and Dermatologists perfectly illustrates how diagnostic uncertainty modulates human-computer trust. While psychological literature suggests a general trend toward "algorithm appreciation" among laypeople (Logg, Minson, & Moore, 2019), our clinical data reveals a more nuanced reality governed by expertise. For GPs, the AI acts as a highly confident, authoritative agent that artificially resolves the psychological discomfort of uncertainty. By deferring to the machine, GPs engage in cognitive offloading. For experts, however, the opaque AI is perceived as an epistemic threat rather than a collaborative tool, preventing them from benefiting from AI as a reliable "safety net" for their own cognitive blind spots.

10.2. Defensive Medicine and Systemic Triage Inversion

The fact that GPs accepted AI-generated false positives 81% of the time provides empirical validation for the systemic threat of "triage inversion" (discussed in Chapter 7). If this behavior scales across a national healthcare system, primary care ceases to function as a gatekeeper. Driven by algorithmic defensive medicine, GPs will flood specialized dermatology departments with algorithmically flagged benign lesions. This will inevitably paralyze referral pathways, paradoxically increasing wait times and delaying surgical interventions for patients with genuine, high-risk melanomas.

10.3. The Amplification of Epistemic Injustice

The erratic behavioral responses observed when evaluating lesions on darker skin (Fitzpatrick IV-VI) underscore the clinical danger of algorithmic bias. The AI's inability to confidently assist GPs in diagnosing patients of color means that these populations are structurally excluded from the benefits of "algorithmic triage." If GPs cannot rely on the AI for darker skin tones, and their own baseline morphological training on these skin types is lacking, the healthcare system automates and amplifies existing demographic health disparities.

10.4. Limitations

This study has several limitations. The *in silico* "reader study" design lacks the rich contextual and tactile information of an in-person physical examination (e.g., patient history, palpation of the lesion). The controlled environment may also induce a Hawthorne effect, where participants pay more attention to the AI than they would in a rushed, 10-minute primary care consultation. Furthermore, while the threat of legal liability was frequently discussed, participants did not face real-world legal consequences for their decisions in this simulation. Future research must evaluate human-AI interaction in prospective, real-world clinical trials to validate the magnitude of these behavioral shifts.

11. Conclusions

The integration of Artificial Intelligence into primary care dermatology represents far more than a mere technological upgrade; it is a profound socio-technical disruption that fundamentally rewrites the cognitive, ethical, and systemic rules of clinical practice. While deep learning algorithms continue to achieve unprecedented statistical accuracy in retrospective computational vacuums, this study explicitly demonstrates that a highly accurate algorithm does not seamlessly translate into a safe or effective clinical tool. Medicine is an inherently human endeavor, and the efficacy of medical AI is entirely dictated by the psychological and ergonomic realities of the physician interacting with it.

Our findings reveal a dangerous dichotomy in trust calibration, driven by varying levels of clinical expertise and diagnostic uncertainty. For General Practitioners operating on the frontline, the AI acts as an authoritative oracle that rapidly resolves the discomfort of ambiguity. However, this dynamic breeds profound automation bias, transforming an intended "digital second opinion" into an unquestioned diagnostic mandate. Fueled by the fear of legal liability, GPs engage in algorithmic defensive medicine, blindly accepting false-positive alerts. If deployed unchecked, this behavior threatens to trigger systemic "triage inversion," paralyzing specialized dermatology departments with algorithmically generated false alarms and delaying care for patients with genuine malignancies.

Conversely, for expert dermatologists, the epistemological opacity of black-box algorithms triggers severe algorithmic aversion. Unable to interrogate the machine's hidden reasoning, experts perceive the AI as an epistemic threat to their refined heuristic workflows, rejecting the technology entirely and nullifying its potential as a cognitive safety net.

Furthermore, the uncritical adoption of current AI models poses existential threats to healthcare equity and medical education. The pervasive lack of diversity in AI training data automates epistemic injustice, leaving patients with darker skin phototypes vulnerable to systematic misdiagnosis. Simultaneously, the routine cognitive offloading of diagnostic tasks to machines risks stalling the clinical maturation of young physicians, leading to widespread cognitive atrophy and medical deskilling. A healthcare system overly dependent on an algorithmic crutch—operated by physicians who have lost the ability to independently diagnose—is inherently fragile and fundamentally unsafe.

To harness the genuine life-saving potential of AI in dermato-oncology, the medical, legal, and computer science communities must collectively abandon the myopic pursuit of isolated algorithmic perfection (maximizing the AUC). The future of AI in medicine mandates a radical shift toward Human-Centered Artificial Intelligence (HCAI). This paradigm requires the rigorous implementation of Explainable AI (XAI) interfaces, such as saliency maps and content-based image retrieval, to replace black boxes with transparent, collaborative clinical tools. It requires diverse, globally representative training datasets to ensure algorithmic equity. Finally, it demands a transformation in medical education and liability frameworks to protect physicians from becoming the "moral crumple zones" of flawed technologies.

Ultimately, artificial intelligence should not be designed to replace human clinical judgment, nor to exploit human cognitive vulnerabilities. By acknowledging the limits of the human brain under pressure and designing for calibrated trust, we can forge a true clinician-AI symbiosis—ensuring that the technology serves as a powerful, synergistic safeguard for the patient, rather than a hidden cognitive liability for the physician.

Funding: This research received no external funding.

Conflicts of Interest: The author declares no conflict of interest.

Author Contributions: The author confirms sole responsibility for the following: study conception and design, data collection, analysis and interpretation of results, and manuscript preparation.

REFERENCES

1. Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. <https://doi.org/10.1037/xge0000033>
2. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115–118. <https://doi.org/10.1038/nature21056>
3. Liu, Y., Jain, A., Eng, C., Way, D. H., Lee, K., Bui, P., ... & Webster, D. R. (2020). A deep learning system for differential diagnosis of skin diseases. *Nature Medicine*, 26(6), 900–908. <https://doi.org/10.1038/s41591-020-0842-3>
4. Lowell, B. A., Froelich, C. W., Federman, D. G., & Kirsner, R. S. (2001). Dermatology in primary care: Prevalence and patient disposition. *Journal of the American Academy of Dermatology*, 45(2), 250–255. <https://doi.org/10.1067/mjd.2001.114598>
5. Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
6. Rajpurkar, P., Chen, E., Banerjee, O., & Topol, E. J. (2022). AI in health and medicine. *Nature Medicine*, 28(1), 31–38. <https://doi.org/10.1038/s41591-021-01614-0>
7. Creswell, J. W., & Plano Clark, V. L. (2018). *Designing and conducting mixed methods research* (3rd ed.). SAGE Publications.
8. Jian, J. Y., Bisantz, A. M., & Drury, C. G. (2000). Foundations for an empirically determined scale of trust in automated systems. *International Journal of Cognitive Ergonomics*, 4(1), 53–71. https://doi.org/10.1207/S15327566IJCE0401_04
9. Lyell, D., & Coiera, E. (2017). Automation bias and verification complexity: A systematic review. *Journal of the American Medical Informatics Association*, 24(2), 423–431. <https://doi.org/10.1093/jamia/ocw105>
10. World Medical Association. (2013). World Medical Association Declaration of Helsinki: Ethical principles for medical research involving human subjects. *JAMA*, 310(20), 2191–2194. <https://doi.org/10.1001/jama.2013.281053>
11. Blease, C., Kaptchuk, T. J., Bernstein, M. H., Mandl, K. D., Halamka, J. D., & DesRoches, C. M. (2019). Artificial intelligence and the future of primary care: Exploratory qualitative study of UK general practitioners' views. *Journal of Medical Internet Research*, 21(3), e12802. <https://doi.org/10.2196/12802>
12. Croskerry, P. (2009). A universal model of diagnostic reasoning. *Academic Medicine*, 84(8), 1022–1028. <https://doi.org/10.1097/ACM.0b013e3181ace703>
13. Goddard, K., Roudsari, A., & Wyatt, J. C. (2012). Automation bias: A systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association*, 19(1), 121–127. <https://doi.org/10.1136/amiajnl-2011-000089>
14. London, A. J. (2019). Artificial intelligence and black-box medical decisions: Accuracy versus explainability. *Hastings Center Report*, 49(1), 15–21. <https://doi.org/10.1002/hast.973>
15. Elish, M. C. (2019). Moral crumple zones: Cautionary tales in human-robot interaction. *Engaging Science, Technology, and Society*, 5, 40–60. <https://doi.org/10.17351/ests2019.260>
16. Froomkin, A. M., Kerr, I., & Pineau, J. (2019). When AIs outperform doctors: Confronting the challenges of a tort-induced over-reliance on machine learning. *Arizona Law Review*, 61, 33–99.
17. Maliha, G., Gerke, S., Cohen, I. G., & Parikh, R. B. (2021). Artificial intelligence and liability in medicine: Balancing safety and innovation. *The Milbank Quarterly*, 99(3), 629–647. <https://doi.org/10.1111/1468-0009.12504>
18. Price, W. N., Gerke, S., & Cohen, I. G. (2019). Potential liability for physicians using artificial intelligence. *JAMA*, 322(18), 1765–1766. <https://doi.org/10.1001/jama.2019.15064>
19. Adamson, A. S., & Smith, A. (2018). Machine learning and health care disparities in dermatology. *JAMA Dermatology*, 154(11), 1247–1248. <https://doi.org/10.1001/jamadermatol.2018.2348>
20. Daneshjou, R., Smith, M. P., Sun, M. D., Rotemberg, V., & Zou, J. (2021). Lack of transparency and potential bias in artificial intelligence data sets and algorithms: A scoping review. *JAMA Dermatology*, 157(11), 1362–1369. <https://doi.org/10.1001/jamadermatol.2021.3129>

21. Zou, J., & Schiebinger, L. (2018). AI can be sexist and racist—it's time to make it fair. *Nature*, 559(7714), 324–326. <https://doi.org/10.1038/d41586-018-05707-8>
22. Asan, O., Bayrak, A. E., & Choudhury, A. (2020). Artificial intelligence and human trust in healthcare: Focus on clinicians. *Journal of Medical Internet Research*, 22(6), e15154. <https://doi.org/10.2196/15154>
23. Esmailzadeh, P. (2020). Use of AI-based tools for healthcare purposes: A survey study from consumers' perspectives. *BMC Medical Informatics and Decision Making*, 20(1), 170. <https://doi.org/10.1186/s12911-020-01191-1>
24. Kerasidou, A. (2020). Artificial intelligence and the ongoing need for empathy, compassion and trust in healthcare. *Bulletin of the World Health Organization*, 98(4), 245–250. <https://doi.org/10.2471/BLT.19.237198>
25. Ferrero, N. A., Morrell, D. S., & Burkhart, C. N. (2020). Skin scan: A demonstration of the need for FDA oversight of direct-to-consumer diagnostic smartphone apps. *Journal of the American Academy of Dermatology*, 82(2), 445–447. <https://doi.org/10.1016/j.jaad.2019.08.064>
26. Gomolin, A., Netchiporouk, E., Gniadecki, R., & Litvinov, I. V. (2020). Artificial intelligence applications in dermatology: Where do we stand?. *Frontiers in Medicine*, 7, 100. <https://doi.org/10.3389/fmed.2020.00100>
27. Nelson, C. A., & Roberson, M. L. (2022). The impact of teledermatology and artificial intelligence on the dermatology workforce and access to care. *Dermatologic Clinics*, 40(1), 101–110.
28. Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8)
29. Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing machine learning in health care—Addressing ethical challenges. *The New England Journal of Medicine*, 378(11), 981–983. <https://doi.org/10.1056/NEJMp1714229>
30. Verghese, A., Shah, N. H., & Harrington, R. A. (2018). What this computer needs is a physician: Humanism and artificial intelligence. *JAMA*, 319(1), 19–20. <https://doi.org/10.1001/jama.2017.19198>
31. Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
32. Sinz, C., Tschandl, P., Rosendahl, C., Akay, B. N., Argenziano, G., Cabo, H., ... & Zalaudek, I. (2017). Accuracy of dermatoscopy for the diagnosis of nonpigmented cancers of the skin. *Journal of the American Academy of Dermatology*, 77(6), 1100–1109. <https://doi.org/10.1016/j.jaad.2017.07.022>