LESION
LENS

AI-powered clinical decision support for skin lesions

1

Problem Space

Global Impact of Melanoma

Vocabulary

Malignant: Cancerous
Benign: Non-Cancerous
Lesion: Mole

Melanoma Incidence and Mortality

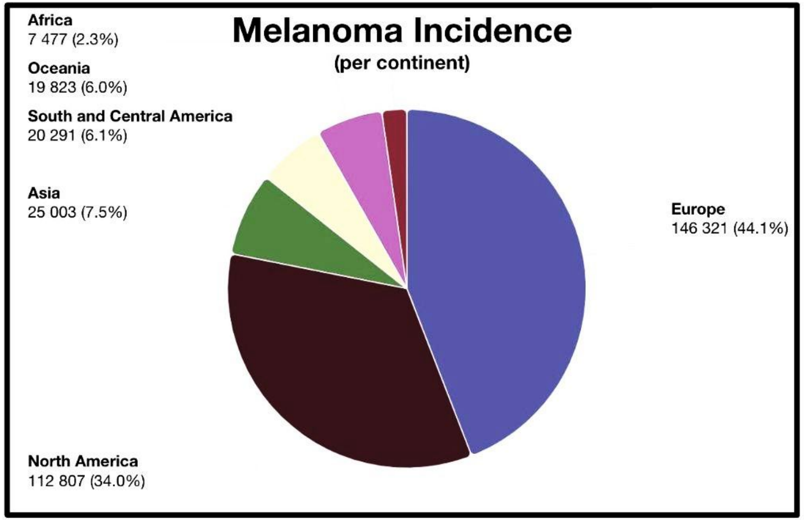

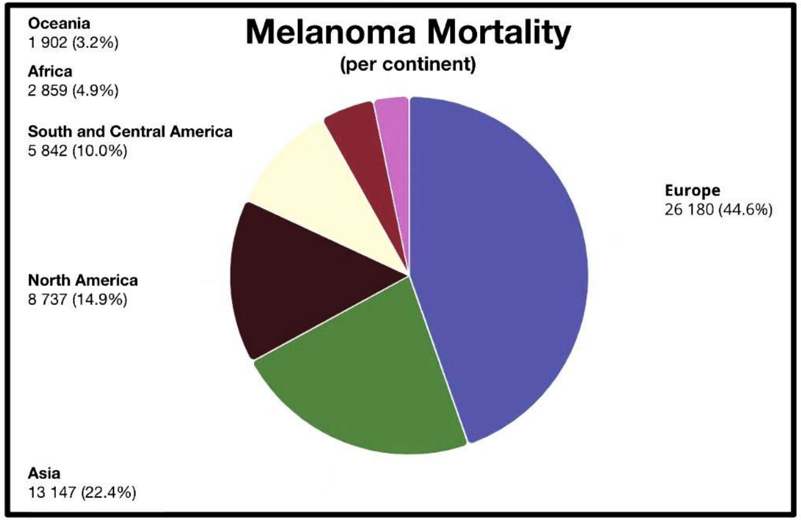

Barriers and Risks

Barriers to Treatment

- Difficult visual assessment
- Inefficient diagnosis process
  - **99.9%** unnecessary biopsy
- Cost: **~$150**
- Wait time: **~78 days**

Health Risks

stage of skin cancer	5-year survival rate
0	99-100%
1	99.4%
2	82-94%
3	32-93%
4	29.8%

2

Our Solution

Application Workflow

User takes photo → **User ensures focus and centeredness** → **User uploads photo to application** → **Photo is processed and sent to LesionLens model** → **Application provides clinical decision support**

Our Application Demo

Lesion Lens Model Details The App The Team

The Problem

Melanoma is a malignant skin lesion with a substantial global impact, with incidence across the continents ranging from approximately 8,000 in Africa to 150,000 in Europe (World Health Organization, 2020). Early-stage melanomas, along with other malignant skin lesions, are often difficult to differentiate from benign lesions by visual examination alone, even for trained clinicians. Thus, diagnostic processes for skin lesions are inefficient, with unnecessary biopsy rates reaching up to 99.9% (International Skin Imaging Collaboration n.d.), , leading to higher healthcare costs and potentially exacerbating patient anxiety. Additionally, even in areas with access to dermatological care, patients face barriers such as high costs and long wait times. In the United States, the average out-of-pocket cost for an uninsured dermatology appointment is over $150, and wait times for consultations can extend up to 78 days (Walk-in Dermatology,n.d-a, n.d.-b). The health risks associated with delayed or missed melanoma diagnoses emphasize the need to improve visual diagnostic accuracy for skin lesions. The five-year survival rate of this skin cancer ranges from 99% for stage 1 to 30% for stage 4 melanoma (HealthlineHealthline , 2023). Given the rapid growth of smartphones and mobile health solutions, there is a promising opportunity to enhance access to care through direct-to-consumer tele-consultation services for skin lesions. However, challenges remain, particularly regarding the inconsistent quality of patient-acquired images in telemedicine (International Skin Imaging Collaboration, n.d.). Thus, we are motivated to develop a low-cost, non-invasive, and efficient diagnostic tool for skin lesions.

Our Solution

Our solution, LesionLens, is a minimum viable product for a web application that offers near-instant visual diagnostic support, providing insights on potential malignancy risk through a machine learning-powered model that takes in user images. Recognizing the limitations of visual assessment, we aim to provide affordable and rapid support to individuals concerned about skin lesions, with a focus on the following target users: high-risk individuals

Our Offering

Target Users

High-risk individuals with atypical lesions

Healthcare-avoidant individuals

Quick information seekers

Value Proposition

- Clinical decision support
- Affordability
- Direct-to-patient accessibility
- Robustness to quality inconsistencies

3

Technical Approach

3.1 The Data

Data Source

Datasets

1. BCN_20000 Dataset (Department of Dermatology, Hospital Clínic de Barcelona)
2. HAM10000 Dataset (Department of Dermatology, Medical University of Vienna)
3. MSK Dataset (Anonymous)
4. SIIM-ISIC 2020 Challenge Dataset (International Skin Imaging Collaboration)

Combined Raw Data

- Combined 4 datasets, de-duplicated and filtered
- N = 20,000 images
- Target = benign vs. malignant



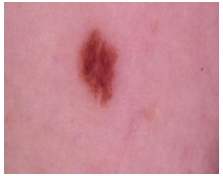


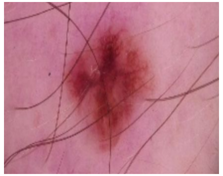

Exploratory Data Analysis

1. Color in Image Data

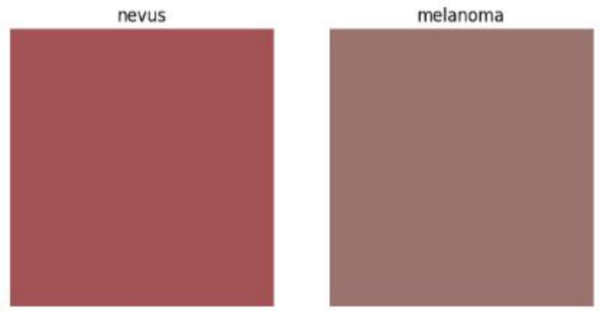

demonstrates importance of color in images.

3. Target Variable Imbalance

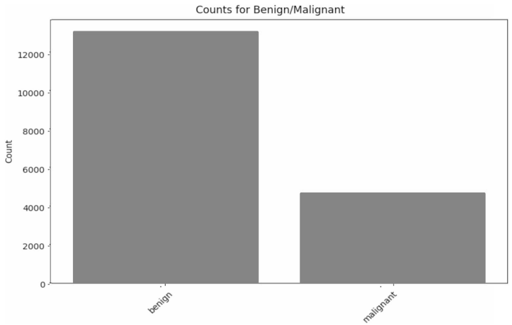

highly imbalanced classes.

2. Categorical Variables in Metadata

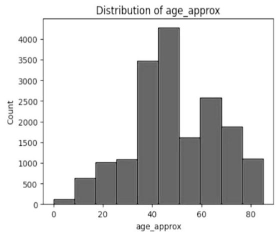

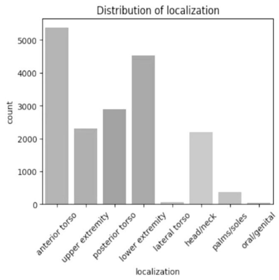

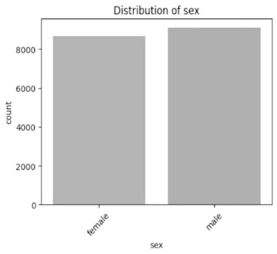

didn’t include metadata in model due to limited value.

broad age range, common torso localization, even sex split.

Dataset Preparation

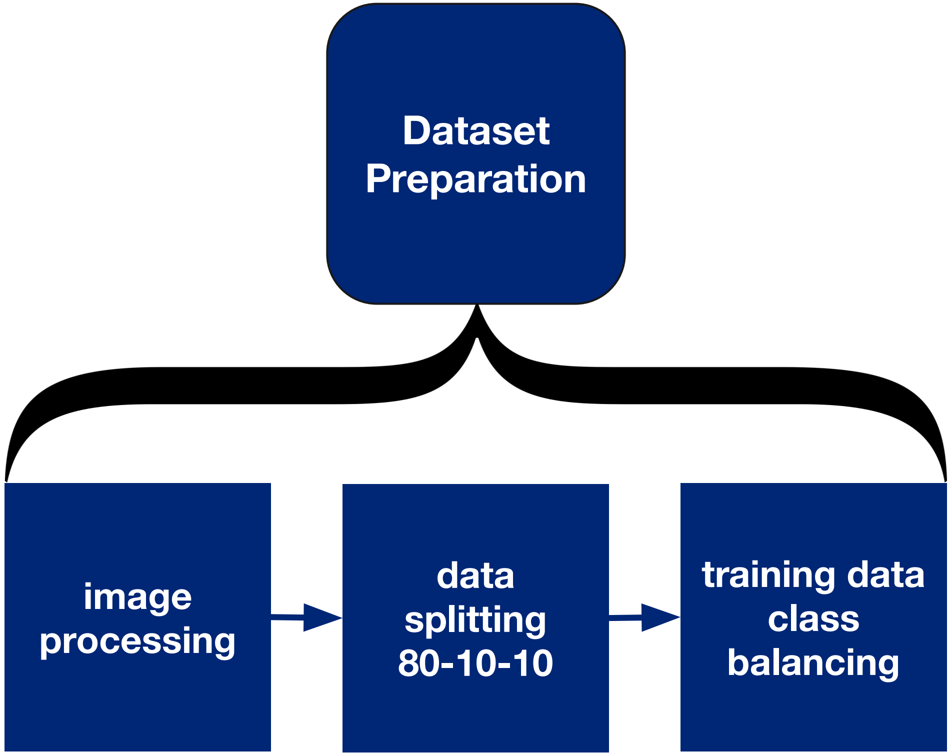

Image Processing

all raw data undergoes image processing, which consists of 4 key stages.

resizing → **hair removal** → **masking** → **normalizing**

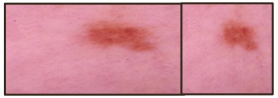

to make dimensions compatible with vision models.

to remove hair, noise, and other artifacts from images.

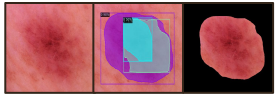

to isolate lesion from extraneous features in user images.

to bring RGB pixel color into a standard range compatible with vision models.

Note: Our application applies this image processing to user image before feeding it into the classification model. Hair removal and masking ensure stability across diverse user images.

Data Splitting

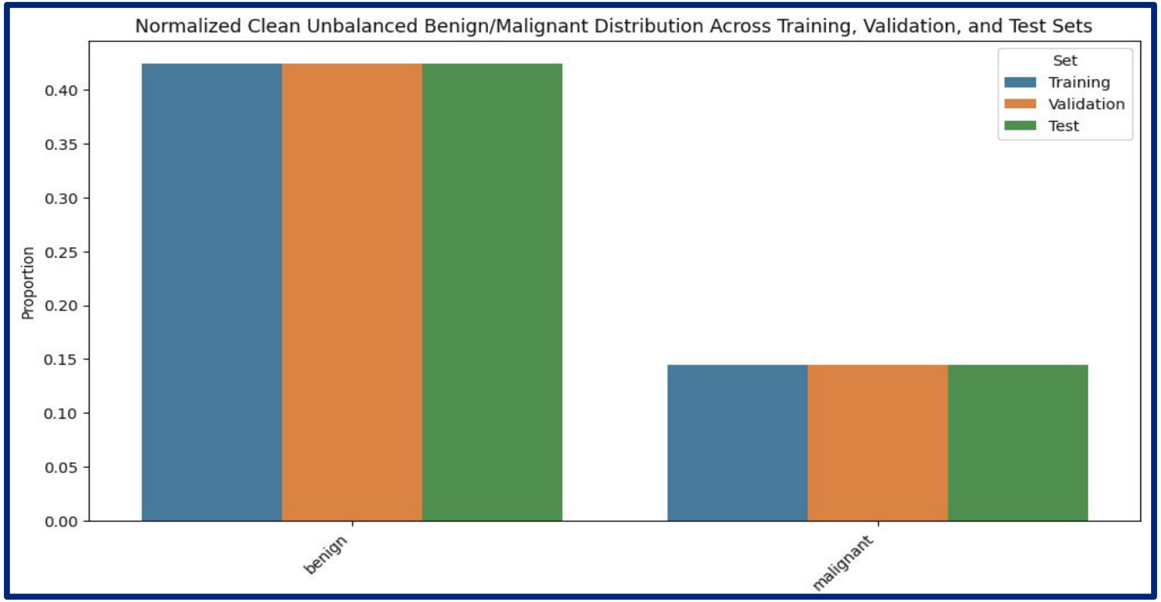

stratified split 80-10-10 into **training**, **validation**, and **test** sets preserves class distribution for more reliable performance metrics and improved generalizability.

Class Balancing

Class Distribution Before and After

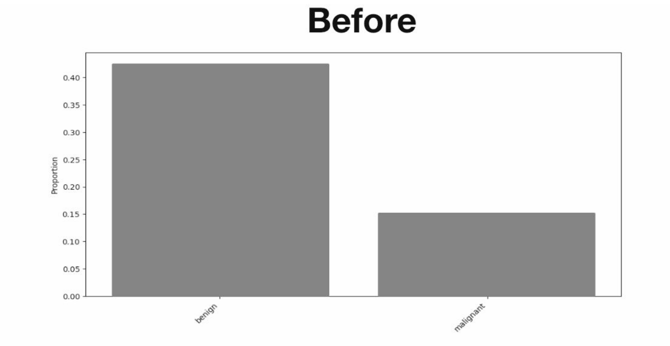

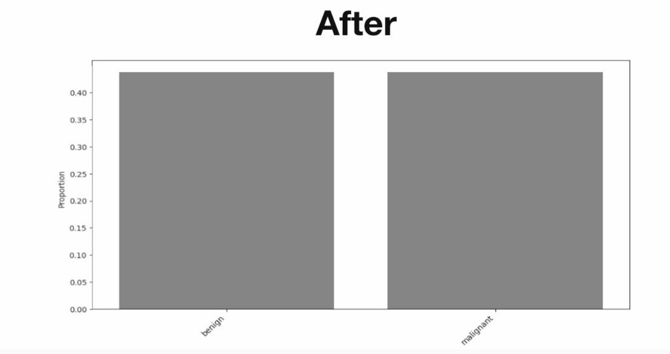

Layered Augmentation

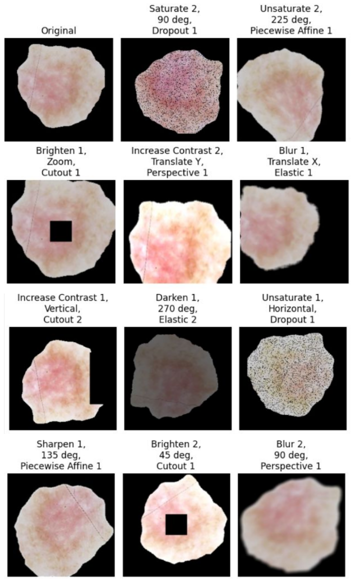

Benefits of Augmentation

1. Better predictions for underrepresented class.
2. Robustness to noise and variation in image quality.

3.2 Modeling

Modeling Approach

Approach

- Tuned **custom CNNs** and **transfer learning models.**
- Intended to select "best" subset of models for ensemble.
- Balancing **precision**, **recall**, **accuracy**, **F1 score**, and **AUC**.

Design Choices

- Tested **lightweight** (e.g., MobileNet) and **dense** (e.g., SENet) models.
- Tuned for goals like high **recall**, **speed**, or addressing **misclassifications.**

Training Techniques

- Optimized with **learning rate schedules**, **class weights**, **dropout**, and **early stopping**.
- Adjusted settings for memory efficiency.

Model	Precision	Recall	F1 Score	Accuracy	AUC
Custom CNN V1	0.535	0.5563	0.5455	0.7638	0.6955
Custom CNN V2	0.5411	0.6247	0.5799	0.7694	0.7218
SkinLesNet Custom CNN	0.4543	0.6358	0.5299	0.7126	0.6873
Mobilenet V1	0.6691	0.6071	0.6366	0.8234	0.7522
Mobilenet V2	0.6974	0.585	0.6363	0.8296	0.7491
EfficientNetB4	0	0	0	0.7452	0.5
DenseNet121 V1	0.6436	0.6578	0.6507	0.82	0.7667
DenseNet121 V2	0.7861	0.649	0.711	0.8656	0.7943
DenseNet121 V3	0.619	0.695	0.655	0.813	0.774
DenseNet121 V4	0.652	0.669	0.66	0.825	0.773
DenseNet169 V1	0.661	0.65	0.655	0.826	0.768
DenseNet169 V2	0.72	0.545	0.621	0.83	0.736
ResNet50	0.5485	0.574	0.561	0.7711	0.7062
ResNet101	0.5009	0.6203	0.5542	0.7458	0.7045
ResNeXt50	0.4376	0.6115	0.5101	0.7008	0.6714
SEResNeXt101	0.542	0.6976	0.61	0.7728	0.748
SENet154 V1	0.5653	0.6976	0.6245	0.7863	0.7571
SENet154 V2	0.7335	0.5651	0.6384	0.8369	0.7475
SENet154 V3	0.5629	0.6225	0.5912	0.7807	0.7286
InceptionResNet V2	0.674	0.6115	0.6412	0.8256	0.7552
InceptionNet	0.654	0.592	0.621	0.816	0.742
NasNetMobile	0.711	0.532	0.609	0.826	0.73

Model Ensembling

1. **Pruned** from 15 models to 4 models using kappa-error diagram pruning.

Model	Precision	Recall	F1 Score	Accuracy	AUC
Mobilenet V1	0.6691	0.6071	0.6366	0.8234	0.7522
Mobilenet V2	0.6974	0.585	0.6363	0.8296	0.7491
DenseNet121 V1	0.6436	0.6578	0.6507	0.82	0.7667
DenseNet121 V2	0.7861	0.649	0.711	0.8656	0.7943
DenseNet121 V3	0.619	0.695	0.655	0.813	0.774
DenseNet121 V4	0.652	0.669	0.66	0.825	0.773
DenseNet169 V1	0.661	0.65	0.655	0.826	0.768
DenseNet169 V2	0.72	0.545	0.621	0.83	0.736
SEResNeXt101	0.542	0.6976	0.61	0.7728	0.748
SENet154 V1	0.5653	0.6976	0.6245	0.7863	0.7571
SENet154 V2	0.7335	0.5651	0.6384	0.8369	0.7475
SENet154 V3	0.5629	0.6225	0.5912	0.7807	0.7286
InceptionResNet V2	0.674	0.6115	0.6412	0.8256	0.7552
InceptionNet	0.654	0.592	0.621	0.816	0.742
NasNetMobile	0.711	0.532	0.609	0.826	0.73

2. **Ensembled** 4 pruned models using logistic regression stacking.

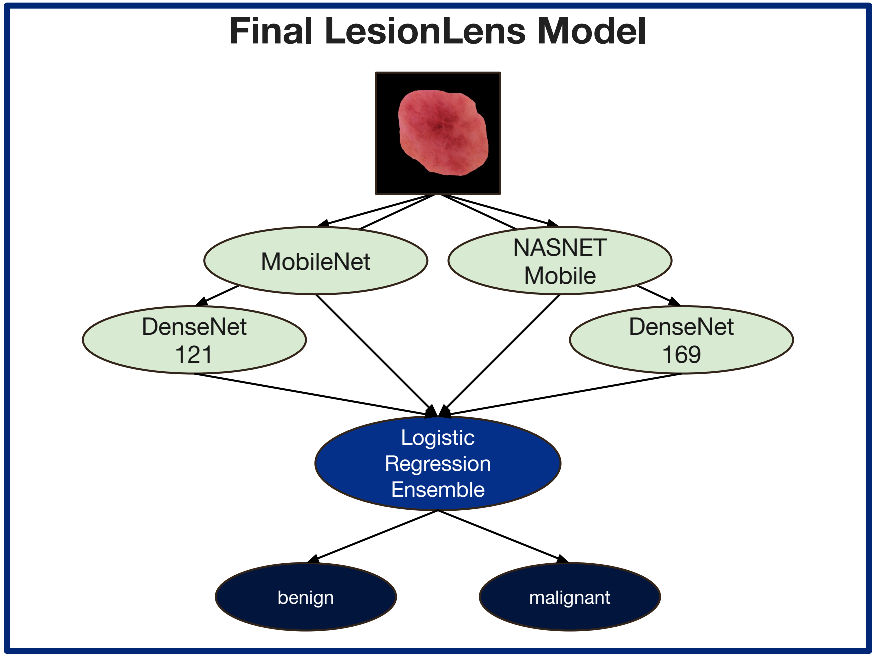

Final Model Benchmarking

Model	Accuracy	Precision	Recall	F1 Score	AUC
*Dermatologist/Human Baseline	84.0%	Not reported.	85.5%	Not reported.	71.0%
**External Best Binary Model	88.8%	Not reported.	83.8%	Not reported.	88.8%
LesionLens Final Model	87.4%	86.7%	87.4%	86.6%	76.1%

*Haenssle, H. A., et al. (2018). Man against machine: Diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists. *Annals of Oncology*, 29(8), 1836–1842. https://doi.org/10.1093/annonc/mdy166

**Pham, T. C., Hoang, V. D., Tran, C. T., Luu, M. S. K., Mai, D. A., et al. (2020). Improving binary skin cancer classification based on best model selection method combined with optimizing fully connected layers of Deep CNN. *2020 International Conference on Multimedia Analysis and Pattern Recognition (MAPR)*, Ha Noi, Vietnam. https://doi.org/10.1109/MAPR49794.2020.9237778

Model Deployment

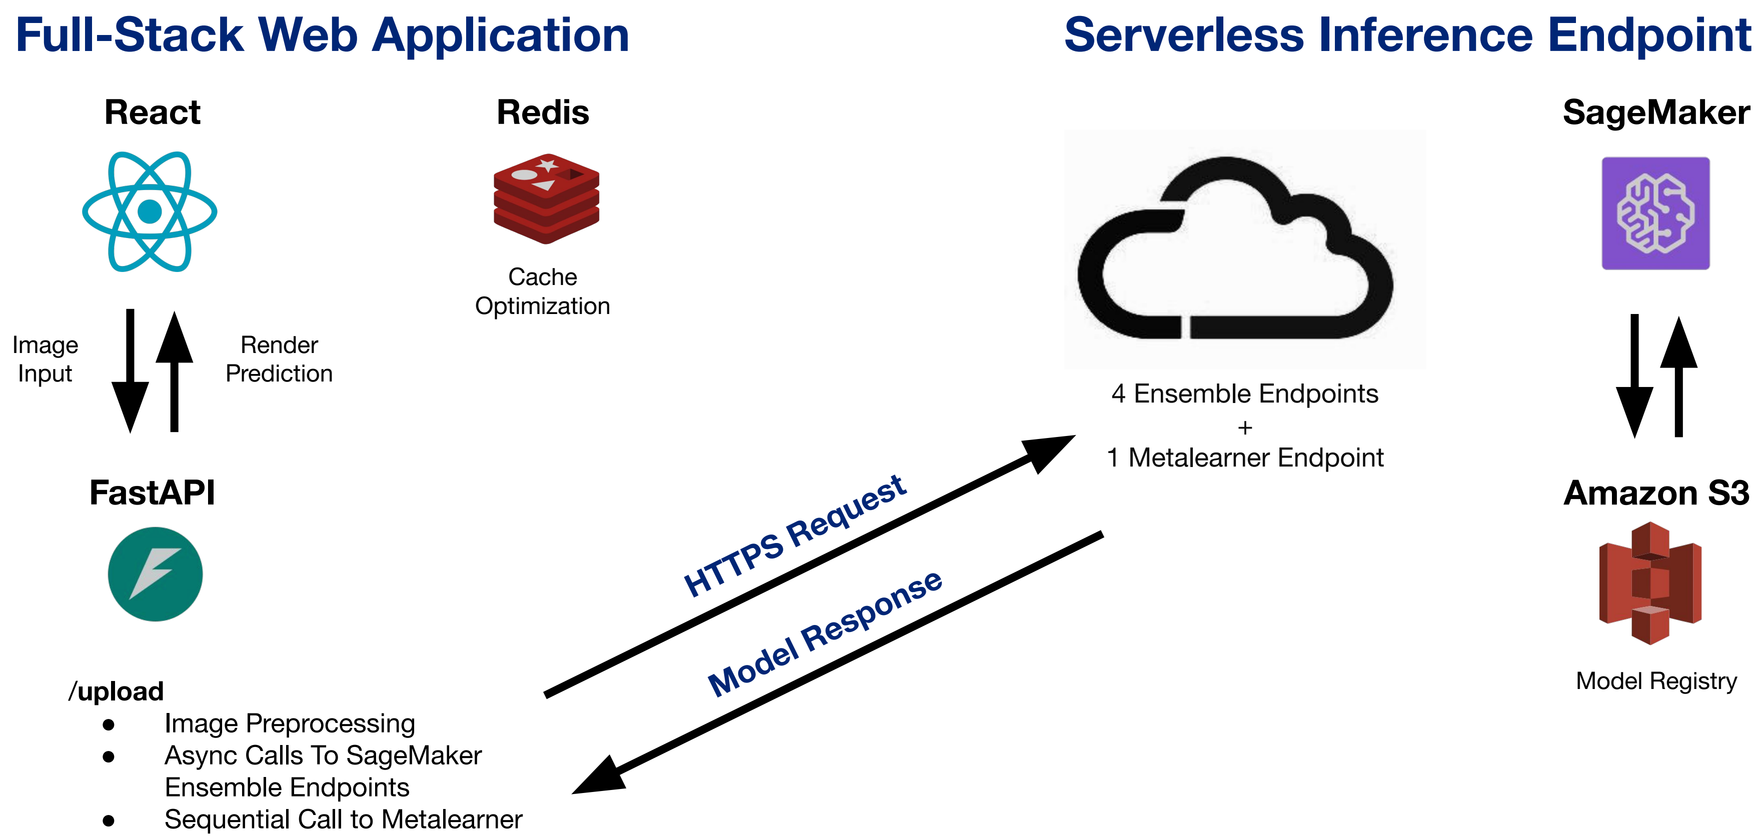

4

Future Work

Future Work

Modeling

Diverse Validated Data

Building a broader set of skin tones, conditions, and imaging, and collaborating with dermatologists for ground-truth validation.

GenAI Capabilities

Incorporating interactive and personalized informational support.

Multiclass Classification Models

Providing decision support for specific diagnosis rather than just malignancy.

Infrastructure

Dockerization

Improving application portability and removing system compatibility issues.

Mobile User-Interface

Improvements informed by usability studies with doctors and patients.

Scaling

Optimizing model for edge devices and enabling offline functionality for community health workers without internet access.

Ethical Considerations

Bias and Fairness

- Ensure robust performance across all skin tones and demographics using diverse datasets.

Patient Safety

- Minimize false negatives and emphasize the tool as assistive, not diagnostic.

Privacy and Security

- Protect user data with encryption, anonymization, and compliance with HIPAA.

Our Mission

LesionLens has the potential to improve early skin cancer detection, alleviate costs, and reduce patient anxiety by providing informative, quick assessments. Our mission is to harness the power of AI to make dermatology care more accessible, equitable, and effective for all.

References

World Health Organization. (2020). Global Health Observatory: Cancer incidence and mortality. Retrieved from https://gco.iarc.who.int/media/globocan/factsheets/cancers/16-melanoma-of-skin-fact-sheet.pdf

International Skin Imaging Collaboration. (n.d.). About ISIC. Retrieved from https://www.isic-archive.com/mission

Walk-in Dermatology. (n.d.-a). How much does it cost to see a dermatologist without insurance? Retrieved from https://walkindermatology.com/how-much-does-it-cost-to-see-a-dermatologist-without-insurance/

Walk-in Dermatology. (n.d.-b). How long do dermatologist referrals take? Retrieved from https://walkindermatology.com/how-long-do-dermatologist-referrals-take/

Healthline. (2023). What Are the Prognosis and Survival Rates for Melanoma by Stage? Retrieved from https://www.healthline.com/health/melanoma-prognosis-and-survival-rates.

Department of Dermatology, Hospital Clínic de Barcelona. (n.d.). BCN_20000 Dataset.

ViDIR Group, Department of Dermatology, Medical University of Vienna. (n.d.). HAM10000 Dataset. https://doi.org/10.1038/sdata.2018.161

Anonymous. (n.d.). MSK Dataset. Retrieved from https://arxiv.org/abs/1710.05006 and https://arxiv.org/abs/1902.03368

International Skin Imaging Collaboration. (2020). SIIM-ISIC 2020 Challenge Dataset. https://doi.org/10.34970/2020-ds01