AvalancheGuard

AI-powered Avalanche Observations

About us

Amiya Ranjan

Ana Zapata

Mandy Korphi

AvalancheGuard was developed as a Capstone project for the Master of Information and Data Science at the School of Information, University of California, Berkeley. We worked under the guidance of our advisors, Joyce Shen and Zona Kostic.

Table of Contents

01
Problem Space

Introductio

Avalanche forecasts save lives. They help people who travel, live, work, and recreate in the backcountry make informed decisions about their activities in snowy mountain environments.

Forecaster Dreams

The executive director of the Sierra Avalanche Center & the former executive Director of AIARE (American Institute of Avalanche Research and Education) want more data.

More data
“We want more people to tell us about avalanche activity they’re seeing”

Timely data
“Especially, if someone triggers an avalanche, we want to know that before the next morning”

The State of Avalanche Forecasting

Snowpack Analysis

Weather Information

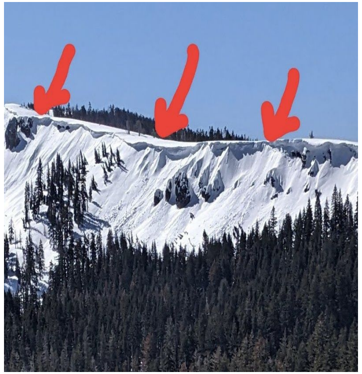

Field Observations

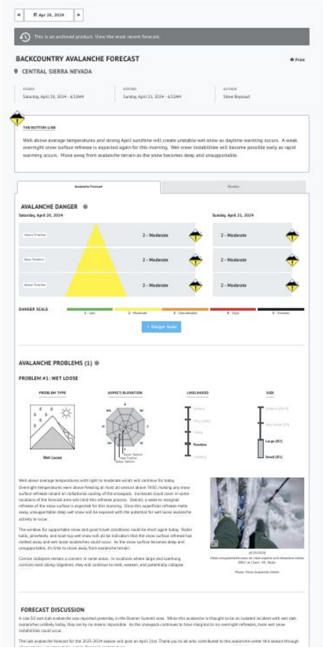

Opportunity to crowdsource for stronger forecasts

Assemble official forecast

Why crowdsource?

Only 3

Official Forecasters in the entire Tahoe Area

141,000

People who read Avalanche Forecast in Tahoe Region

642 (<0.5%)

Observations in the Tahoe area last season.

What Skiers think

Barriers

Recreational skiers don't submit observations because they are nervous about writing the wrong thing.

A desire to help

Skiers would happily upload observations to our application if they knew it was helping the community. They would need guidance and reassurance on what constitutes a useful contribution.

Our Target Users

Image: https://www.cntraveler.com/story/how-to-safely-try-backcountry-skiing

Skiers

Sierra Avalanche Center Facebook Post

Forecasters

Impact

Observations

Number of observations logged.

Accuracy

Observation labeling accuracy of avalanche activity and red flags.

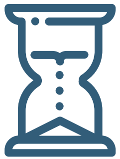

Time

Number of hours saved in an avalanche forecaster typical day

02

Our Solution

Our Solution - Demo

Let us walk through our app from the point of view of our target users.

Image: https://www.cntraveler.com/story/how-to-safely-try-backcountry-skiing

Skiers

Sierra Avalanche Center Facebook Post

Forecasters

MVP Application Architecture

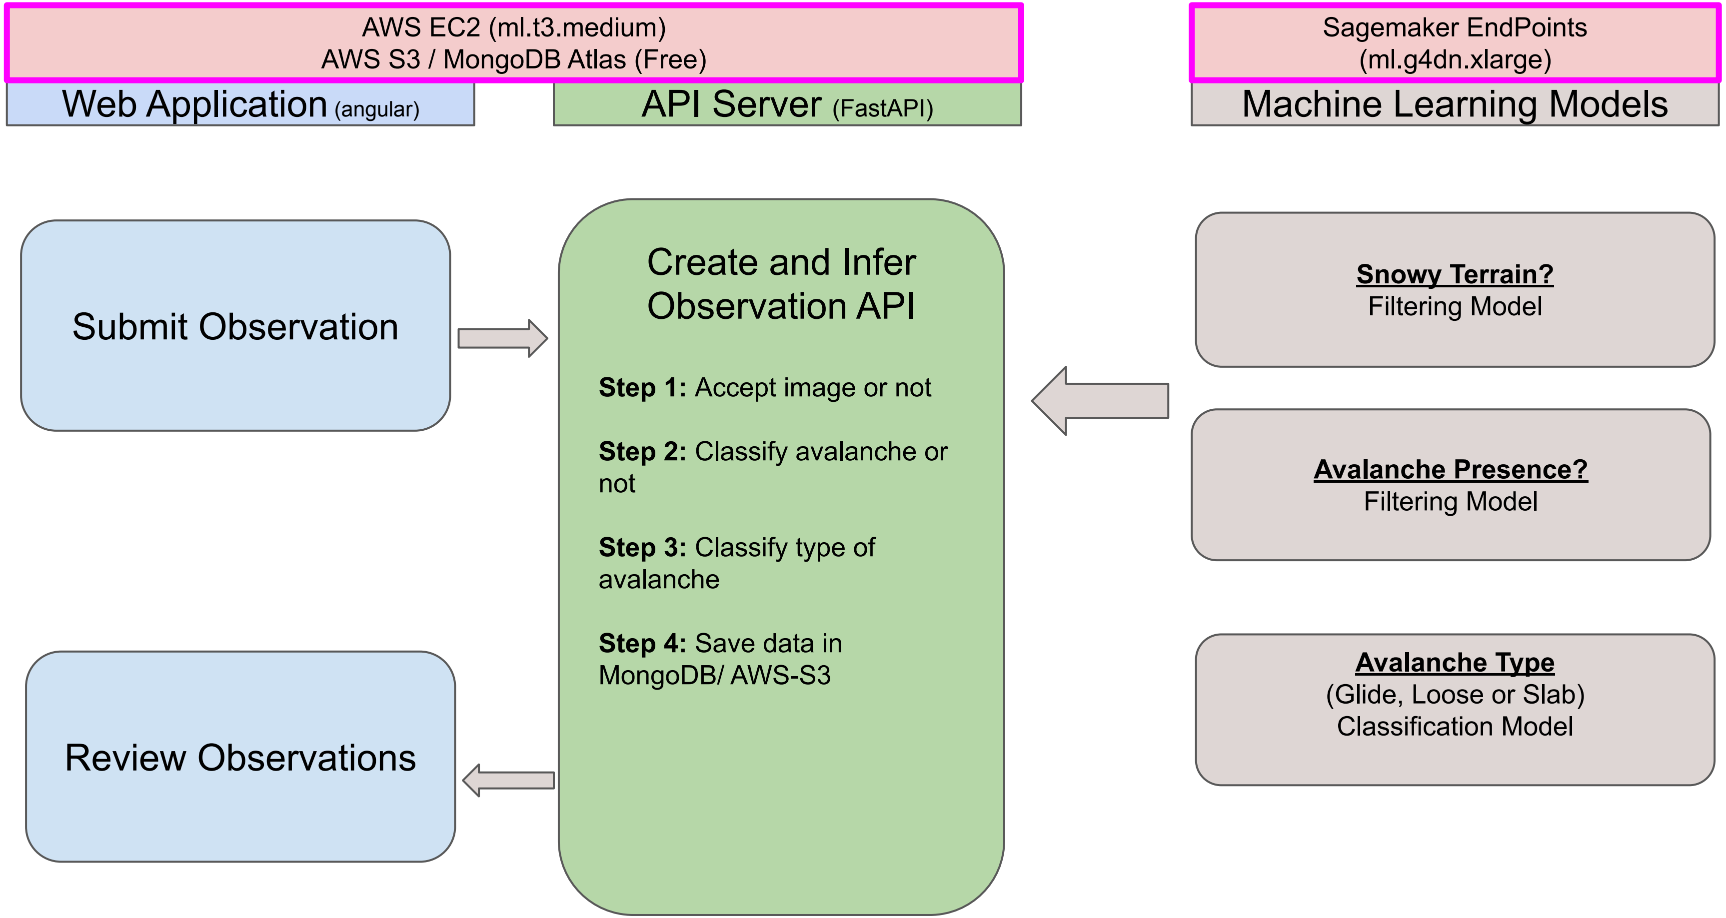

03

Technical Approach

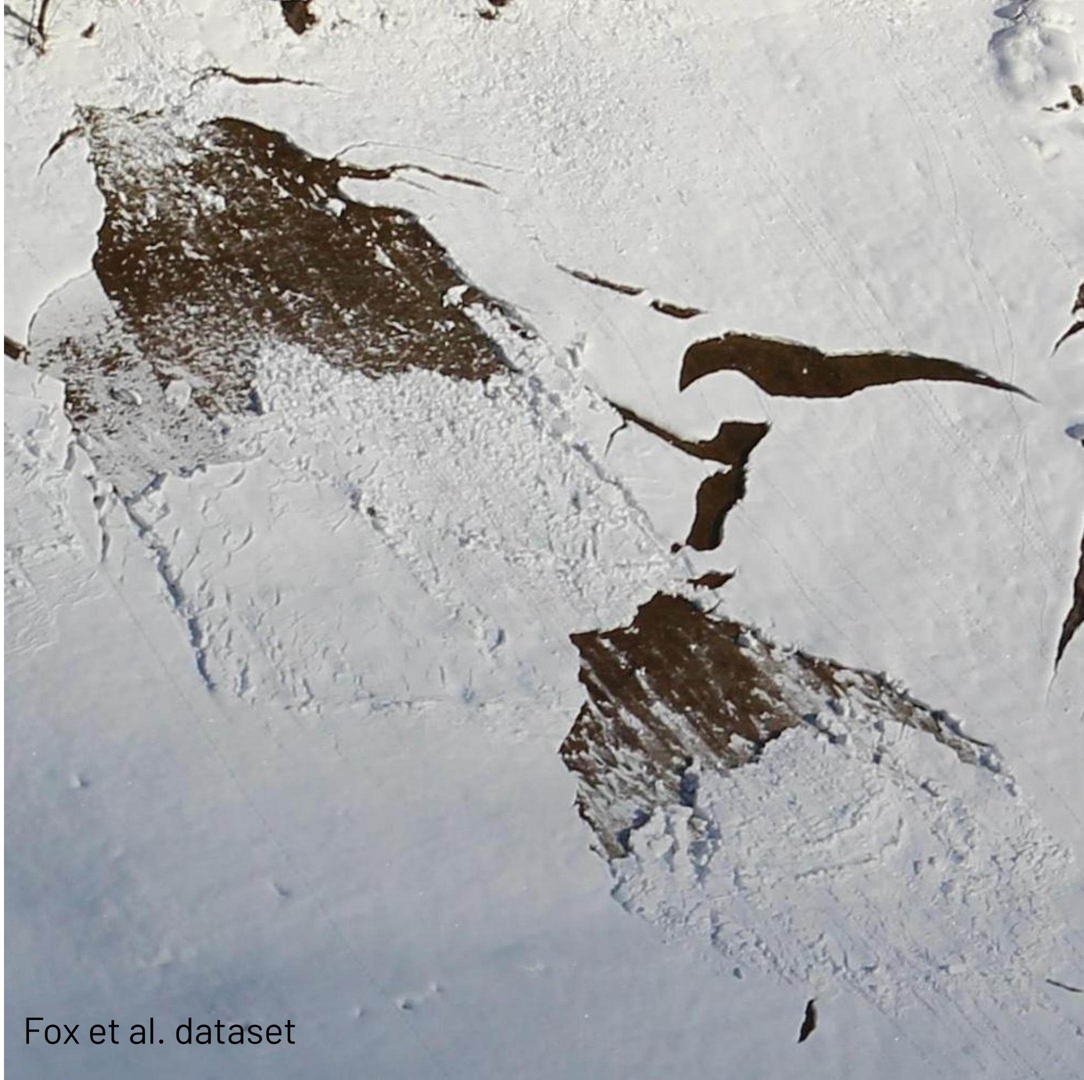

Dataset

Datasets Summary

		Snowy Terrain? Filtering Model	Avalanche Presence? Filtering Model	Avalanche Type Classification Model
Imagenet Variety of other types of pictures Stanford University, Princeton University		✓	x	x
Landscapes Dataset Pictures of coast, desert, forest, mountains & glacier Kaggle Dataset, Published by Utkarsh Saxena		✓	x	x
Primary Dataset: Ground-Based images of avalanches University of Innsbruck, Austria github.com/j-f-ox/avalanche-detection	Wide-Angle Original Version	✓	✓	x
	Close-Up Version (Cropped)	✓	✓	✓

Primary Dataset - EDA

- European bias
- Expertly labeled dataset
- High resolution 3-color jpg images
- Wide-angle pictures
- Conflicting labels (8% images have conflicting labels)

Fox et al. dataset

Primary Dataset Pipeline

- **Crop** close-ups of avalanches from original wide-angle pictures.
- **Resize** all images to 224x224, consistent w/ many computer vision foundation models.
- **Balance Classes** w/ undersampling

Primary Dataset – Cropping Procedure

Cropping:

- Simplifies classification w/ close-up shots instead of wide-angle
- Cleans by removing images with conflicting labels
- Augments dataset ~5x

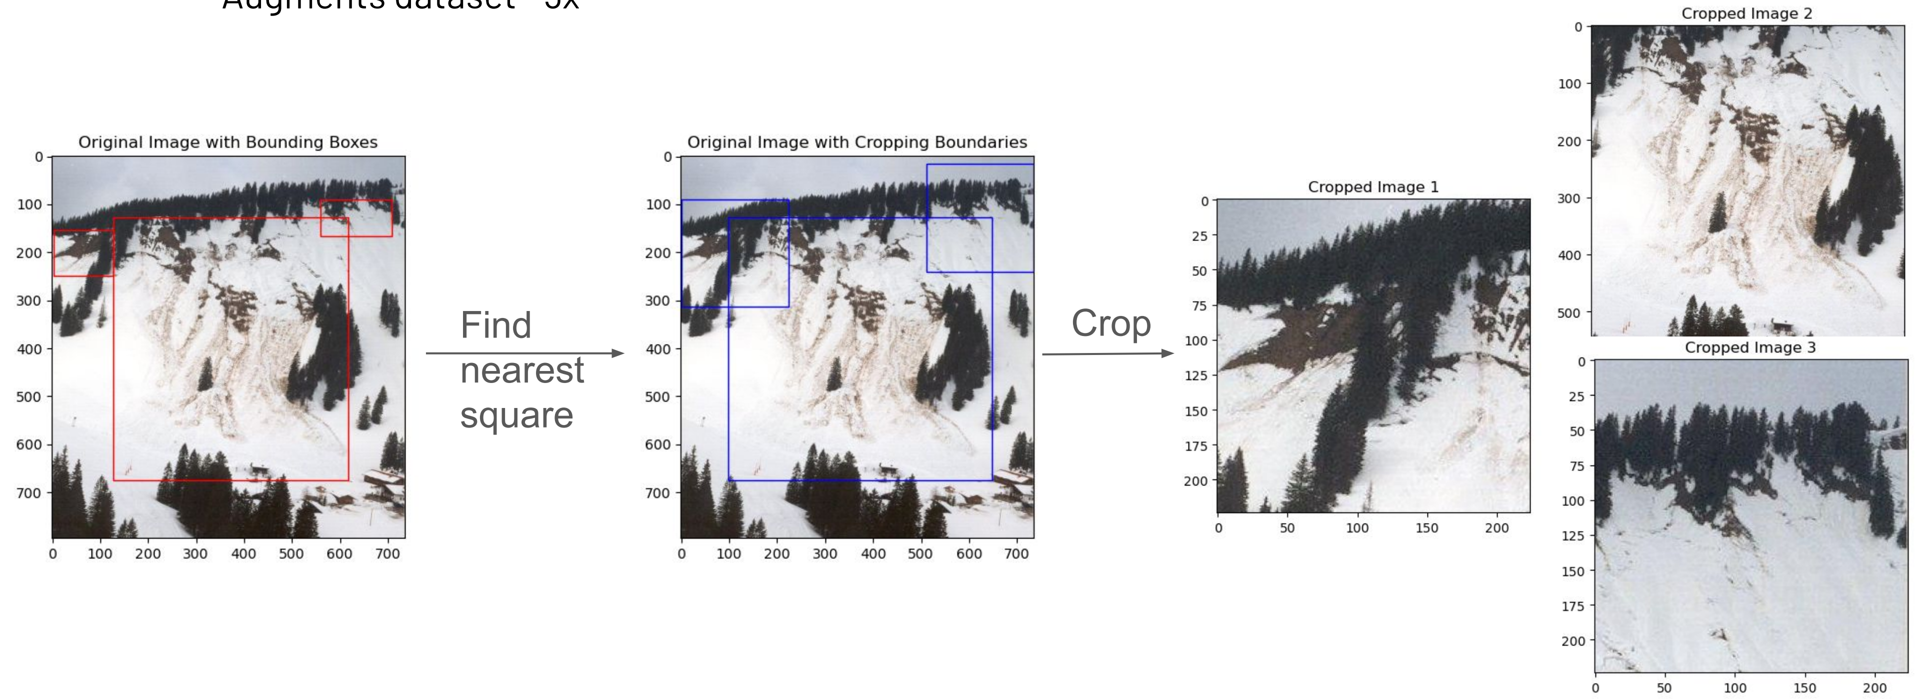

Modeling Approach

Cascaded Model Architecture

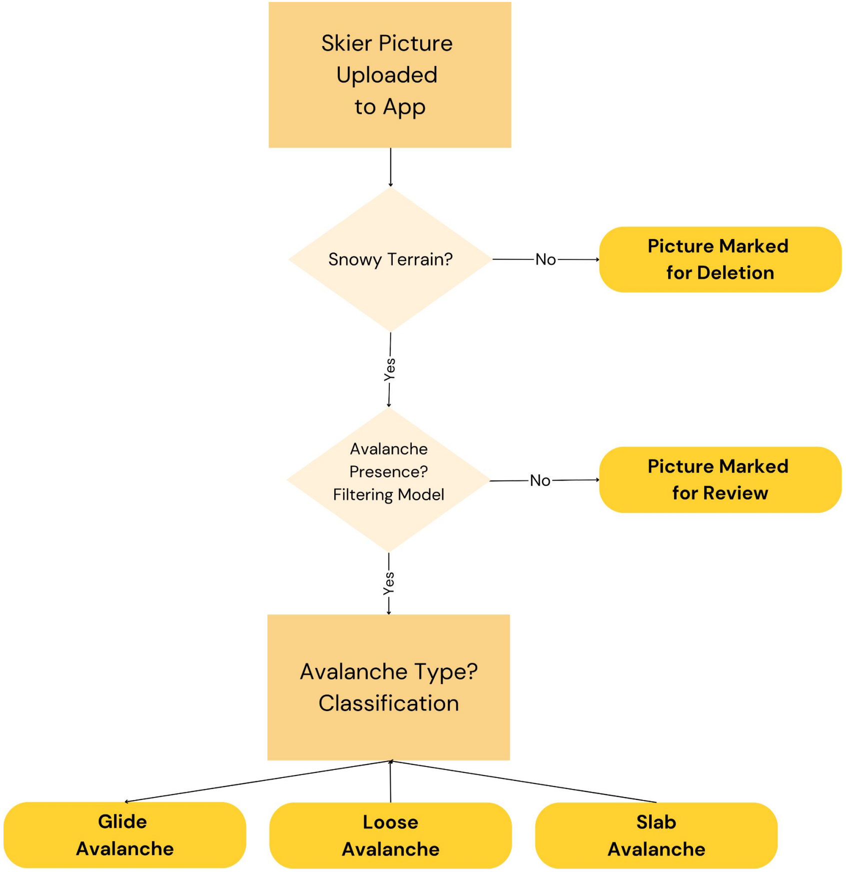

3 Cascaded Models successively take an input picture uploaded by a snow enthusiast and

1. Filter out any pictures that are not a snowy mountain

2. Filter out any pictures that don't feature an avalanche

3. Classify a valid picture of an avalanche into an avalanche type: glide, loose, or slab

Individual Model Architectures

All 3 models are fine-tuned on EfficientNetV2S. Differences in task complexity and nature (binary vs multi-class) drive differences in fine-tuning layers and parameters.

Binary Filtering Models

Input	(224, 224, 3)
Output	Binary
Model	Fine-tune EfficientNetV2S
Finetuning Layers	GlobalAveragePooling2D() Dense(1, activation = 'sigmoid')
Hyperparameters	
Learning Rate	0.001
Optimizer	Adam
Batch Size	32
Epochs	10
Loss Function	binary_crossentropy
Evaluation	accuracy

Avalanche Classification Model

Input	(224, 224, 3)
Output	Glide Avalanche Loose Avalanche Slab Avalanche
Model	Fine-tune EfficientNetV2S
Finetuning Layers	GlobalAveragePooling2D()(x) BatchNormalization()(x) Dropout(0.5)(x) Dense(1024, activation='relu')(x) Dropout(0.5)(x) Dense(num_classes, activation='softmax')(x)
Hyperparameters	
Learning Rate	0.001 to 0.0001 w/ ReduceLROnPlateau & patience 3
Optimizer	Adam
Batch Size	32
Epochs	40 w/ tf.keras.callbacks.ModelCheckpoint
Loss Function	sparse_categorical_crossentropy
Evaluation	accuracy

Model Approach

- We landed on EfficientNetV2S fine-tuning, leveraging a common image classification transfer learning approach for our moderate dataset size.
- Feature extraction of colors and texture was insufficient due to image complexity
- Transformer-based approaches are state-of-the-art but difficult b/c of our small dataset
- Cleaning data, preprocessing, class balancing, fine-tuning layer params were also all tuned

Classification Model Experiment Tracker

Validation Accuracy	Classes	Multiple Labels	Preprocessing	Balancing Approach	Model	Notes
0.3	Glide, Loose, Slide None	yes	equalizeHist GaussianBlur adaptiveThreshold	Undersample	CNN simple	Decisions: 1. Abandon preprocessing 2. Cascade Models
0.4					Resnet Finetune	
0.35					EfficientNetB4	
0.4		removed	none		Color & Texture Features + Logistic Regression	Decision: Color & Texture features insufficient despite small improvements
0.6	Glide, Loose, Slide	Removed by Cropping, effectively augmenting		Augment: - rotate - flip horizontal	Color & Texture Features + Random Forest	
0.56					Color & Texture Features + Logistic Regression	
0.55					Color & Texture TOP 42 Features + Logistic Regression	
0.7					EfficientNetB4	Can't save model! Known TensorFlow Bug
0.73				Undersample		
0.79				Undersample	EfficientNetV2S	Winner

Model Performance

Model Performance – Filtering Models

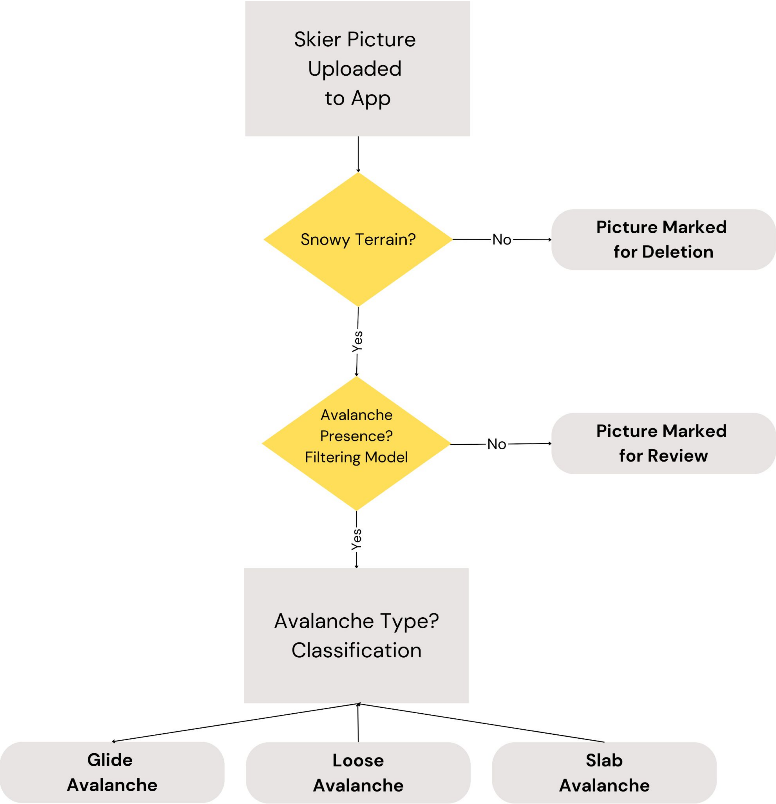

Amazing at filtering out irrelevant pictures
99% accuracy

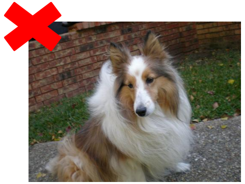

Great at detecting presence of an avalanche
False positives often have skier tracks in the image
84% accuracy

False Positive

Model Performance – Classification Model

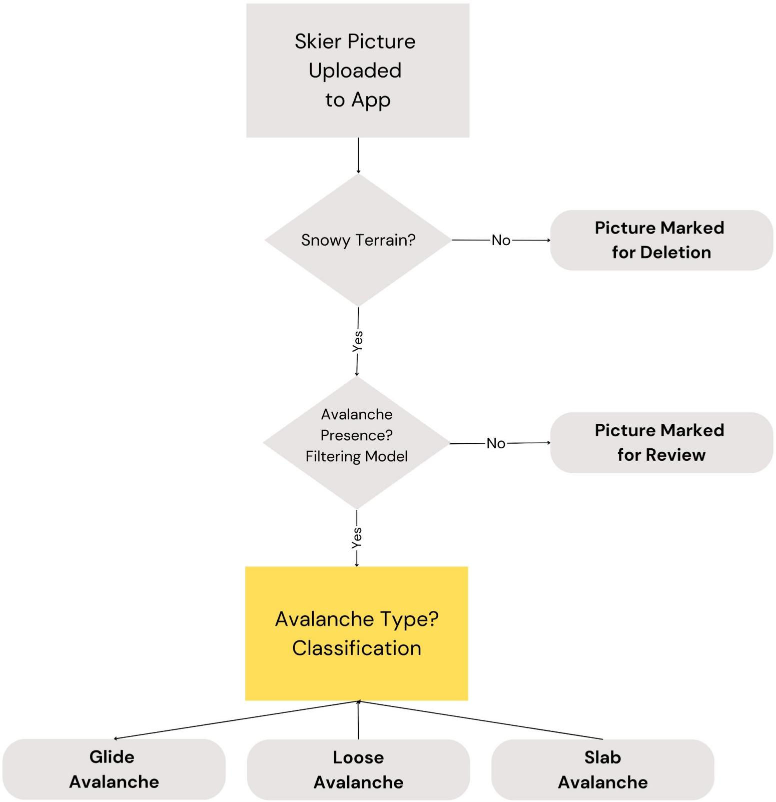

Great at classifying avalanche types for close-up images
79% accuracy

Slab Loose Glide

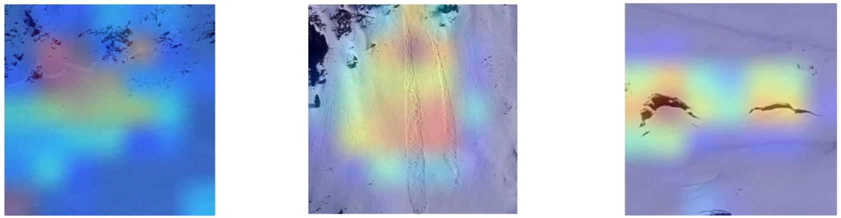

CAM images

Struggles at classifying avalanche types for wide-angle images that include other features such as rocks, trees, clouds
65% accuracy

Sample Complicated Images

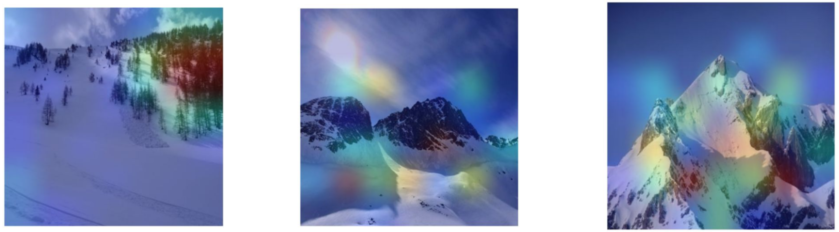

CAM images

Model Performance - Mitigation

For optimal performance with our model, skier pictures should center and focus on the avalanche.

We provide this guidance (and more fun tips!) on our website.

Close-ups more likely anyway from a skier w/ cell phone camera.

Website - Education Section

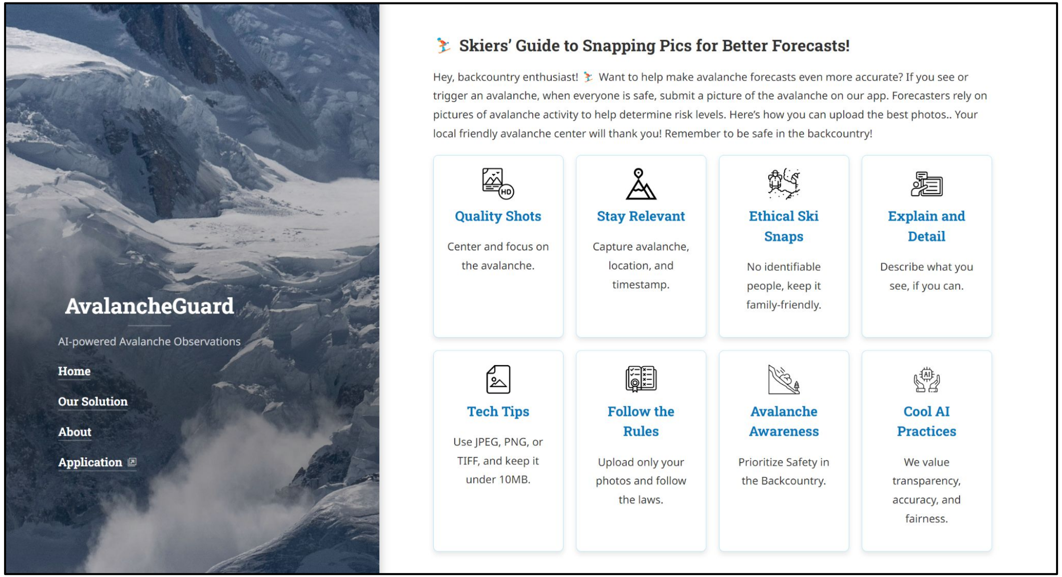

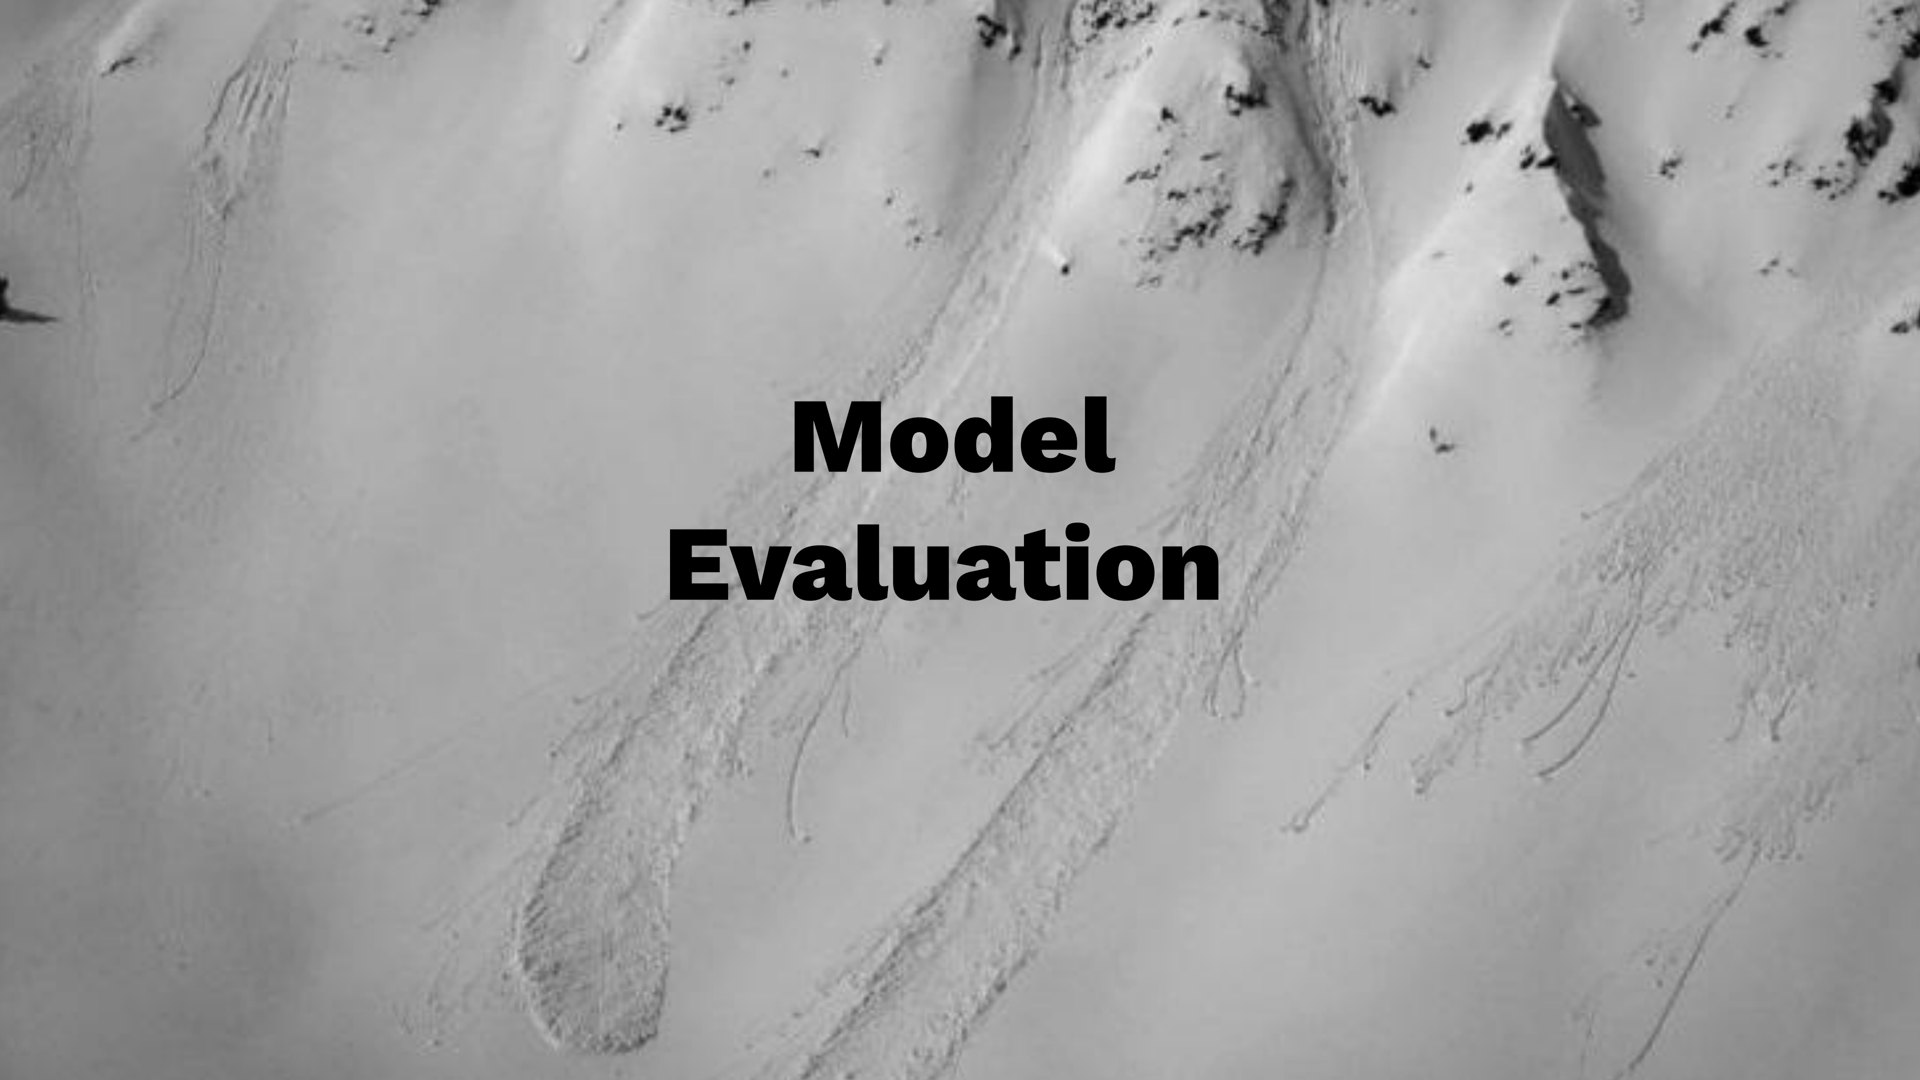

Model Evaluation

Model Evaluation

Models were evaluated using a combination of metrics and analysis tools.
Accuracy was used for model training because the type of misclassification is not important.

Precision, Recall, f1-score

```
Classification Report:
              precision    recall  f1-score   support

       glide       0.86      0.80      0.83       366
       loose       0.78      0.80      0.79       366
        slab       0.74      0.78      0.76       366

    accuracy                           0.79      1098
   macro avg       0.79      0.79      0.79      1098
weighted avg       0.79      0.79      0.79      1098
```

Confusion Matrix

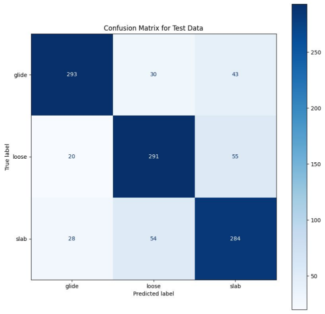

Accuracy and Loss Curves

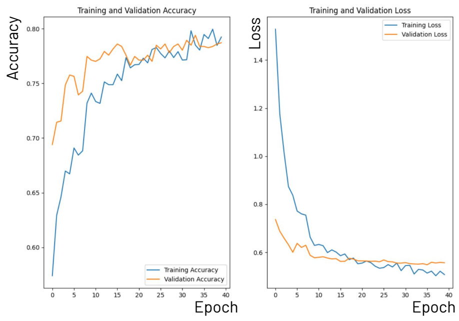

Misclassified Images

Class Activation Maps

Technical Challenges & Takeaways

Architecture Challenge: Single model to detect and classify an avalanche wasn't working

Cascaded models

- Modular
- Future-proof
- Reduce Complexity
- ⏫Explainability

Computer Vision Challenge: Moderately sized dataset

Snow Complexity Challenge:

- Avalanches ≠ Objects
- Lack boundaries, color differences.
- Texture variations subtle
- Rocks, trees, dirt looks similar
- Snow & clouds look similar

EfficientNet Model Selection

- Feature Extraction insufficient
- Transformer-based models needs more data
- EfficientNet (CNN) fine tuning successful for our dataset size & complexity

Closing Remarks

Roadmap

Wide-angle image performance improvements by including a segmentation model

Instagram integration, where many users get their avalanche risk information & already upload images

Expand forecaster-facing observations insight platform to enable sorting and other visualizations

Our Mission

We aim to improve safety in the mountains
by crowdsourcing avalanche observations, and accurately labeling them w/ machine learning,

for avalanche forecasters to obtain
the volume of field data they need to complement their weather and snow pit data

to provide the public with
more accurate and precise avalanche risk forecasts.

Special Thanks

We’d like to thank everyone who advised us and provided feedback throughout our capstone project, especially:

- David Reichel, Executive Director at the Sierra Avalanche Center, for guidance, feedback, and data.
- Richard Bothwell, former Executive Director of AIARE, for guidance, feedback, and avalanche safety insights.
- Jeffrey C, forecaster, for feedback and explaining the forecaster’s process.
- David, Arnaud, Zara, and other backcountry skiers for sharing their experiences and insights.
- Jaeyoung Lim from ASL at ETH Zurich and Elisabeth Hafner from WSL SLF for sharing their avalanche research insights in Switzerland.

Resources

PHOTOS

Slide 5: Images and quotes from Sierra Avalanche Center Facebook posts & avalanche.org

Slide 28: Sierra Avalanche Center Facebook page

Slide 25: Adventure Sports Journal

Adrian Tanic, Unsplash

Paul Bill, Unsplash

Krzysztof Kowalik, Unsplash

Slide 4: Nicolas Cool, Unsplash

Matea Nikolina, Unsplash

Lucas Leonel, Unsplash

Jasper Guy, Unsplash

Slide 16, 43: Alberto Restifo, Unsplash

Slide 2: Anders Jilden, Unsplash

Slide 1, 7, 8, 17: Alessio Soggetti, Unsplash

Slides 16, 20, 24, 28: Fox et al

Information and Links

Slide 6, 7 and 8: Sierra Avalanche Center Annual Report 2020

Datasets

"Automating avalanche detection in ground-based photographs with deep learning". The accompanying dataset of avalanche photographs is available at https://researchdata.uibk.ac.at//records/h07f4-qzd17.

Addison Howard, Eunbyung Park, Wendy Kan. (2018). ImageNet Object Localization Challenge. Kaggle. https://kaggle.com/competitions/imagenet-object-localization-challenge

Landscape Recognition | Image Dataset | 12k Images, https://www.kaggle.com/datasets/utkarshsaxenadn/landscape-recognition-image-dataset-12k-images

Our MVP

Demo 1

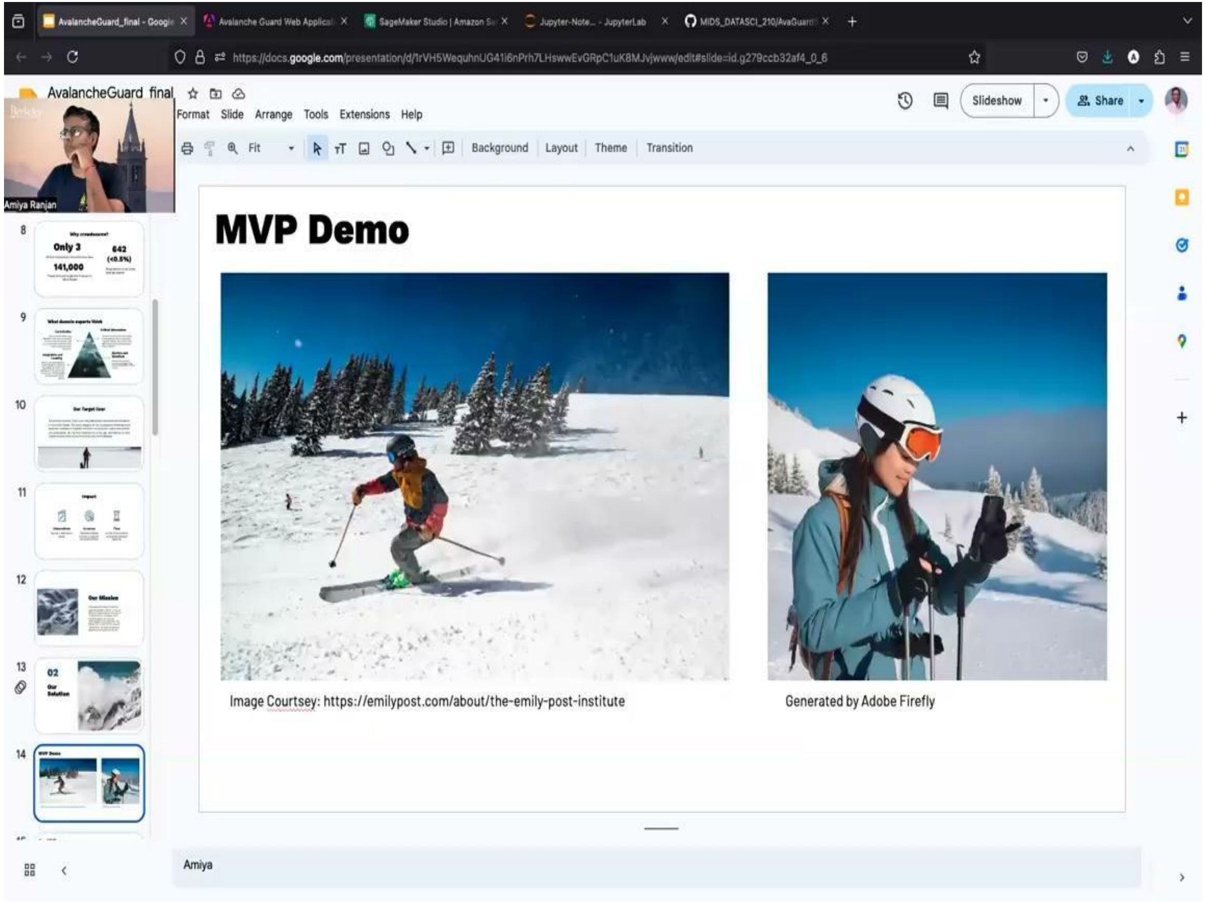

AvalancheGuard website