Adapting hand keypoint detection for leprosy diagnosis

Techniques for adapting the output of existing hand keypoint detectors to predict keypoints for leprosy diagnosis.

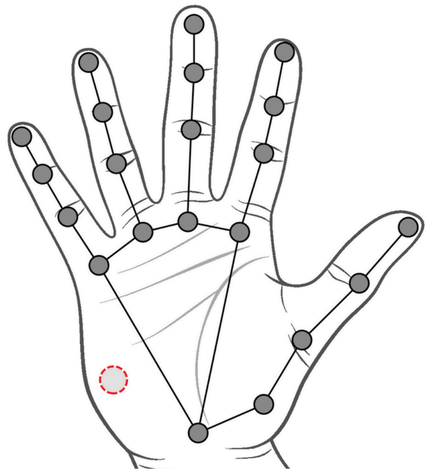

Author
Marek Tran - M.Tran-2@student.tudelft.nl

01 Background

Leprosy, common in developing regions, can be detected early by measuring temperatures at the fingers and the hypothenar point. However, standard hand keypoint detectors do not output the hypothenar. By adapting their outputs to infer this point, temperatures can be automatically measured from infrared images, enabling streamlined leprosy detection.

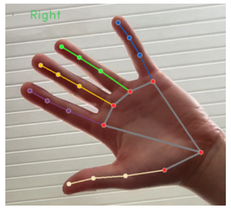

Figure 1: Detected Hand Keypoints by MediaPipe [2]

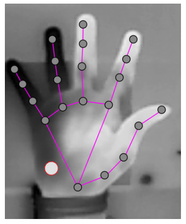

Figure 2: Hand keypoints and Hypothenar Point

02 Research Questions

The main contribution of this study is:

1. Comparison of point prediction models with respect to data availability.
2. RGB/IR dataset of 160 hand images with standard and the hypothenar keypoint manually annotated

The contributions aim answer the following questions:

- **How accurately can a hand keypoint be inferred without data size constraints?**
- **How accurately can the Hypothenar keypoint be inferred with data size constraints?**
- **How do these models generalize to new datasets?**

03 Methodology

The keypoints needed to be normalized and oriented in the same position to focus on the hand proportion and pose and not location within image.

Training models for hypothenar keypoint prediction uses our comparably smaller dataset.

Thumb Keypoint Prediction Experiment

- Lower bound for prediction error when ample data is available (57k images)
- Obtain generalization error between our dataset and external datasets.

Hypothenar Keypoint Prediction Experiment

- Evaluate and compare the models on predicting the hypothenar keypoint (960 images)

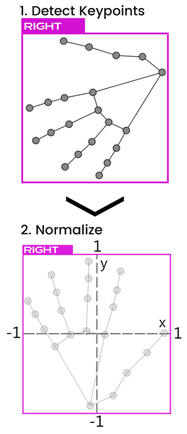

Models compared:

- K-nearest neighbors (K-NN)
- Random Forest
- XGBoost
- Lasso regression
- Schemkes [1] linear interpolation
  - (serving as the baseline)

Mean Euclidean distance in the normalized space was used as loss metric.

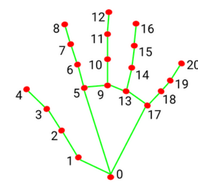

0. WRIST
1. THUMB_CMC
2. THUMB_MCP
3. THUMB_IP
4. THUMB_TIP
5. INDEX_FINGER_MCP
6. INDEX_FINGER_PIP
7. INDEX_FINGER_DIP
8. INDEX_FINGER_TIP
9. MIDDLE_FINGER_MCP
10. MIDDLE_FINGER_PIP
11. MIDDLE_FINGER_DIP
12. MIDDLE_FINGER_TIP
13. RING_FINGER_MCP
14. RING_FINGER_PIP
15. RING_FINGER_DIP
16. RING_FINGER_TIP
17. PINKY_MCP
18. PINKY_PIP
19. PINKY_DIP
20. PINKY_TIP

04 Results

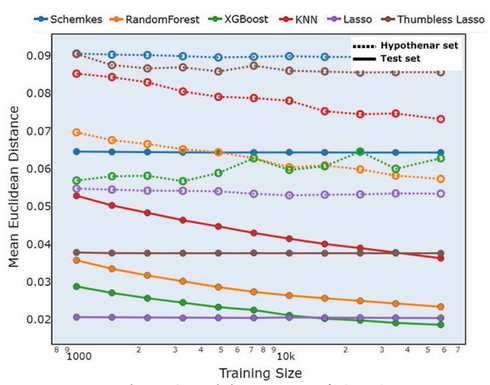

Figure 3: Training Curves of Thumb Knuckle Keypoint Prediction Experiment

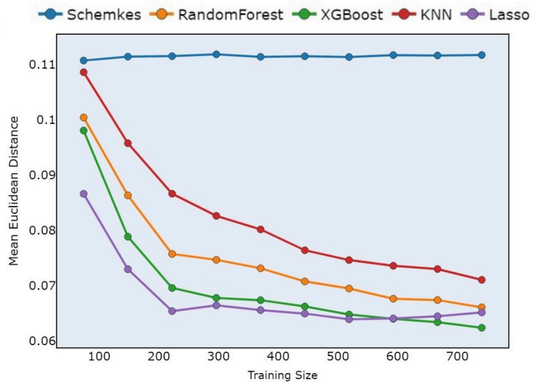

Figure 4: Training Curves of Hypothenar Keypoint Prediction Experiment

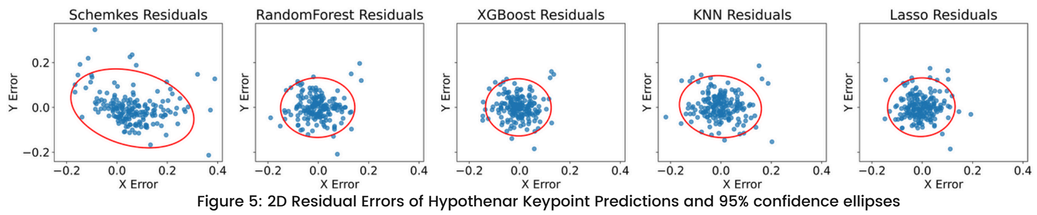

Figure 5: 2D Residual Errors of Hypothenar Keypoint Predictions and 95% confidence ellipses

05 Conclusions

This study highlights the need for research into hand keypoint detection in infrared images It focuses on inferring the hypothenar keypoint, making it possible to use infrared cameras for automated temperature measurements to streamline leprosy detection.

- The best performing model is XGBoost and second best, Lasso Regression.
- More data can reduce the loss of parametric models (XGBoost, K-NN, Random Forest) by up to 35%.
- Lasso regression achieves the lowest loss with low training data volume but is largely dependent on the PINKY_MCP keypoint.
- XGBoost is the overall best performing model but may be prone to overfitting.
- With approximations, XGBoost is expected to predict within 1.19cm and Lasso regression within 1.23cm, with 95% confidence.
- All models outperformed the baseline linear interpolation.

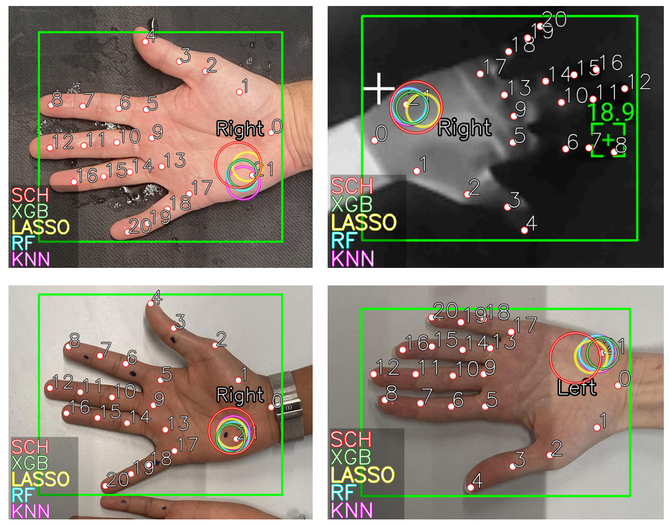

Fig 6. Sample visualisations of the prediction results with 95% confidence circles. Hypothenar Ground-Truth keypoint annotation numbered 21

Supplementary Content
Code available at: https://github.com/MarekTran/hypothenar-inferrence

References

[1] I.A.S. Schemkes. Semi-automatic temperature analysis based on real-time hand landmark tracking in infrared videos. Master's thesis, Delft University of Technology, July 2024
[2] Fan Zhang, Valentin Bazarevsky, Andrey Vakunov, An- drei Tkachenka, George Sung, Chuo-Ling Chang, and Matthias Grundmann. Mediapipe hands: On-device real-time hand tracking, 2020.

Supervisors
Zhi-Yi Lin
Thomas Markhorst
Dr. Jan van Gemert

TU Delft