Automatic Measurement of Patellar Tilt using Deep Learning Methods

Samuel Hunter,[1] Tofunmi Omiye,[2] Marissa Lee[3]

[1]MS Candidate, Computer Science, Stanford University
[2]MS Candidate, Health Policy, Stanford University
[3]PhD candidate, Mechanical Engineering, Stanford University

Stanford Computer Science

Background

Patellar instability

- **What it is:** A clinical syndrome where the patella bone disarticulates from the patellofemoral joint in the knee
- **Who it affects:** Adolescents ages 14-18, with ~148 cases per 100,000
- **How it's measured:** Patellar tilt is the angle between the condyle (lateral to medial) line and the patella (lateral to medial) line
- **How surgery decisions are made:** 10 ± 4.3° is normal, and surgically viable cases are 16 ± 3.3°
- **Why deep learning can help:** Radiologist labels suffer from inter-observer variation

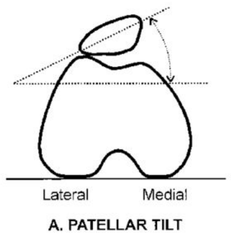

A. PATELLAR TILT

Dataset

- **Source:** JUPITER (JUstifying Patellar Instability Treatment by Early Results) group, a subset from Cincinnati Children's Hospital Medical Center
- **Size:** 216 MRI volumes
- **Dimensions:** ~30 slices per MRI, 256 x 256 to 1024 x 1024 pixels each
- **Split:** 80% train, 20% test, no validation

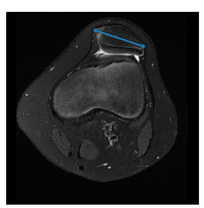
Patella (above) and condyle (below) lines

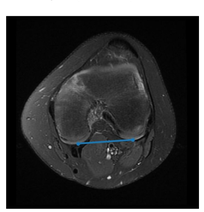

Methods

U-Net

- **What it is:** A fully convolutional architecture with symmetrical contractive and expansive paths
- **Model type**: Segmentation
- **Input:** Any sized 3 channel image
- **Output:** Same size as input, but with k output channels for the k classes
- **Advantages:** Scales to any input, works well with elastic deformations for small datasets

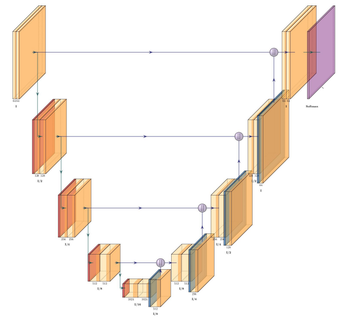

Elastic deformation

- **What it is:** A data augmentation technique to produce new training samples
- **Input:** Any sized 3 channel image
- **Output:** Input perturbed by random vector displacement
- **Advantages:** Low memory overhead, trains model to be invariant to deformed tissue

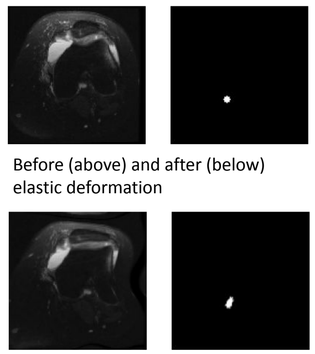
Before (above) and after (below) elastic deformation

Training

- **What:** 4 U-Nets initialized with pretrained weights
- **How:** Adam optimizer, 5 x 10^{-3} learning rate, 32 batch size, 100 epochs
- **Loss:** Cross entropy

Prediction

- **Given:** Prediction mask
- **How:** Softmax over class channels, blur, max localization
- **Output**: x-y keypoint coordinates

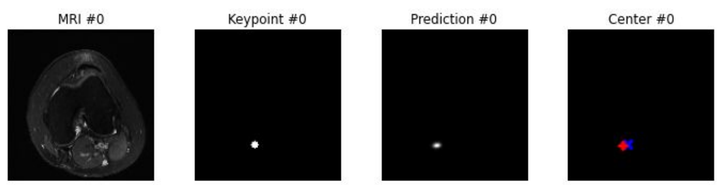

Results

Keypoint coordinates

- **Measure of error:** L2 distance
- **Baseline (modified ResNet):** 20 pixels average error

Keypoint model	Number of Outliers	Error (pixels) Avg.	Max
Lateral condyle	1	3.395	26.653
Medial condyle	3	2.885	27.625
Lateral patella	2	2.445	8.285
Medial patella	1	2.568	15.987

Table 1. Keypoint prediction error statistics.

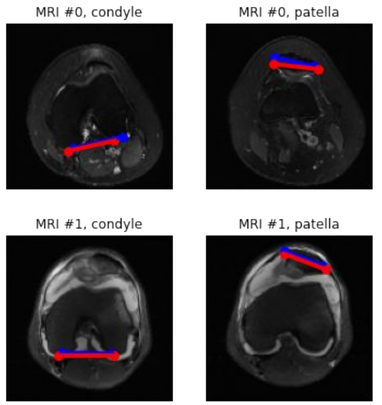

Patellar Tilt

- **Measure of error:** intraclass correlation coefficient (ICC)
- **Baseline:** Human vs. Human

Model	ICC2
U-Net vs. Human	0.555
U-Net (augmented data) vs. Human	**0.768**
Human vs. Human	0.816

Table 2. Patellar tilt reliability results.

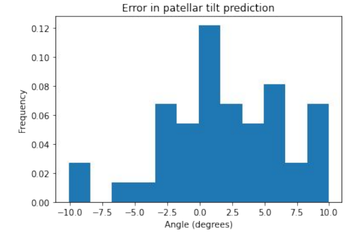

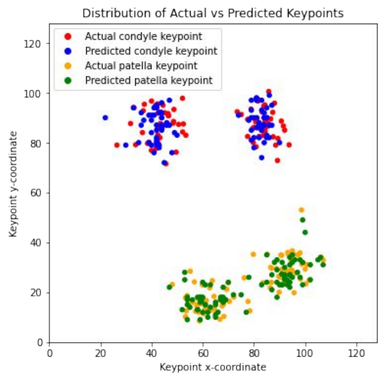

Difficulties

- Multiple or no bright spots in predicted keypoint masks
  - Solution: Train for more epochs
  - Solution: Replace outliers with training median
- Predicted mask won't converged, still looks like input
  - Solution: Relax elastic deformation strength