Face and Pose Tracking

Kat Bradley
Kaylin Spitz

General Layout (Detection)

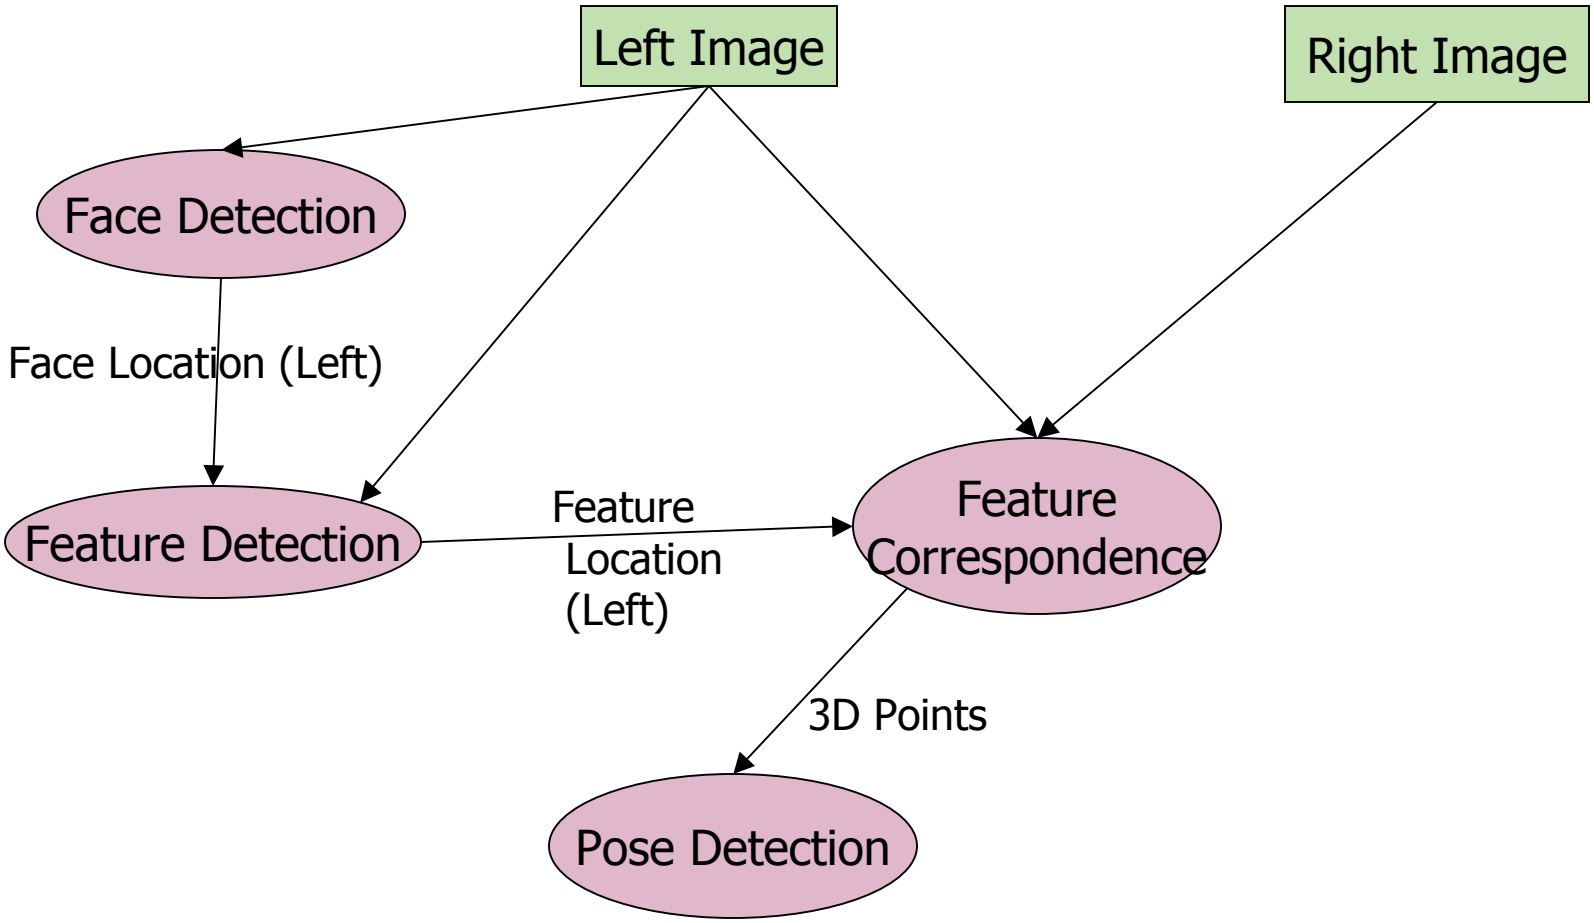

General Layout (Tracking)

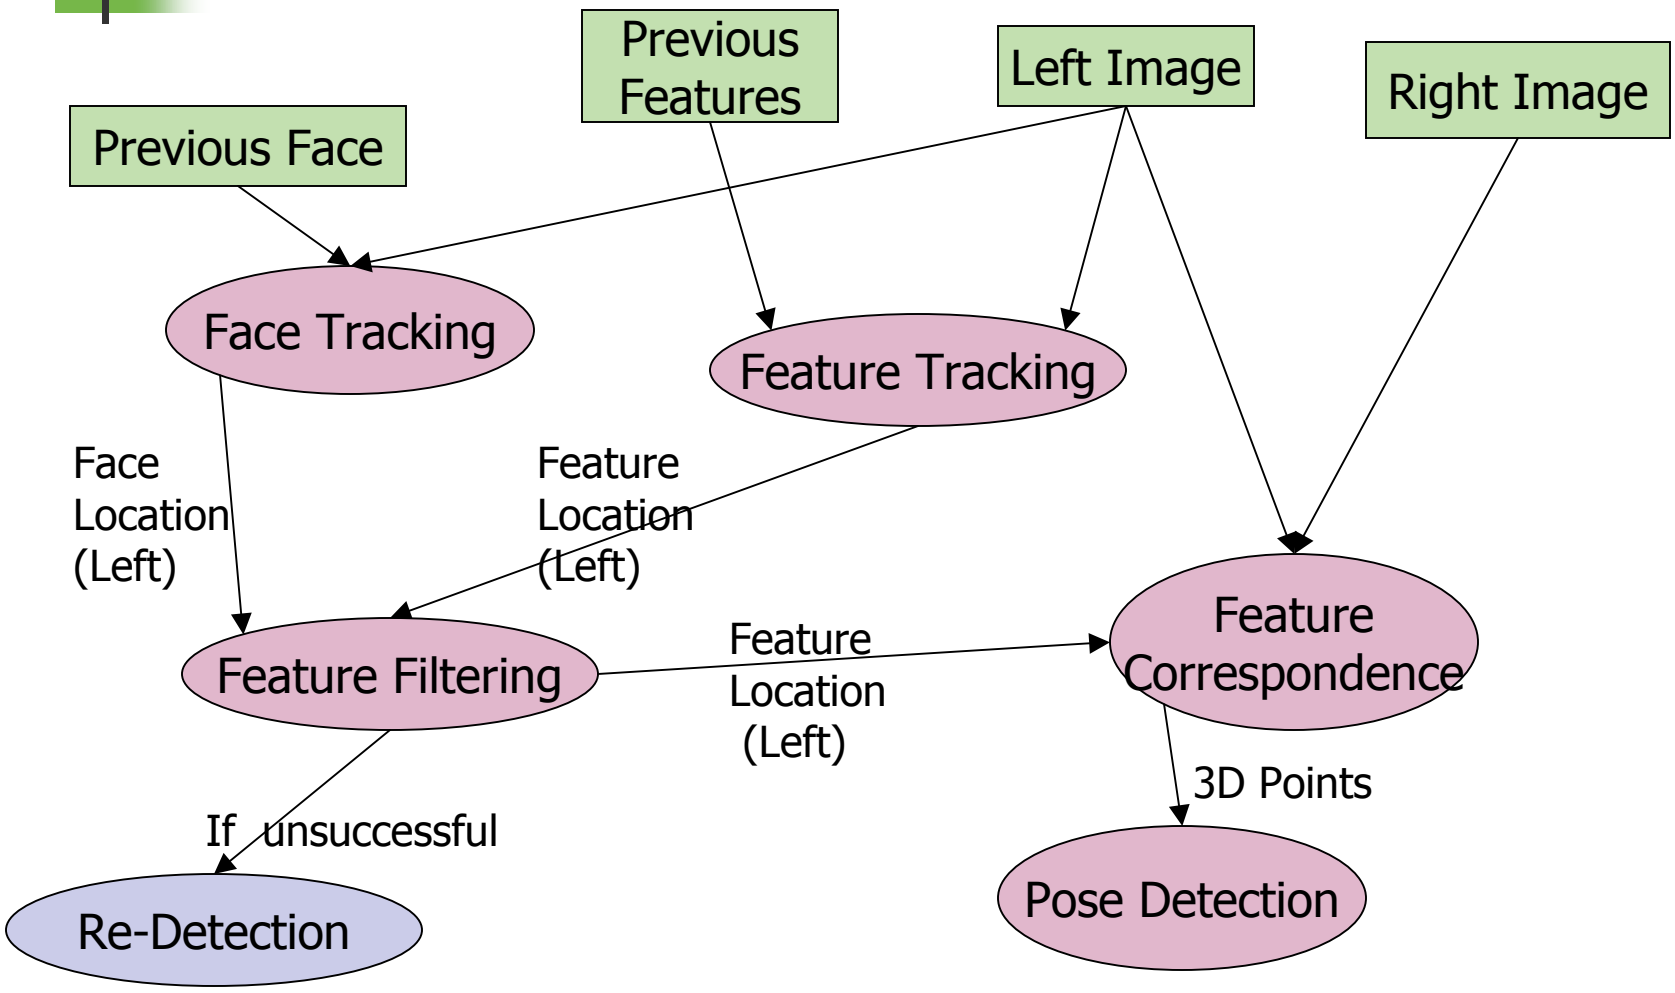

Face Detection

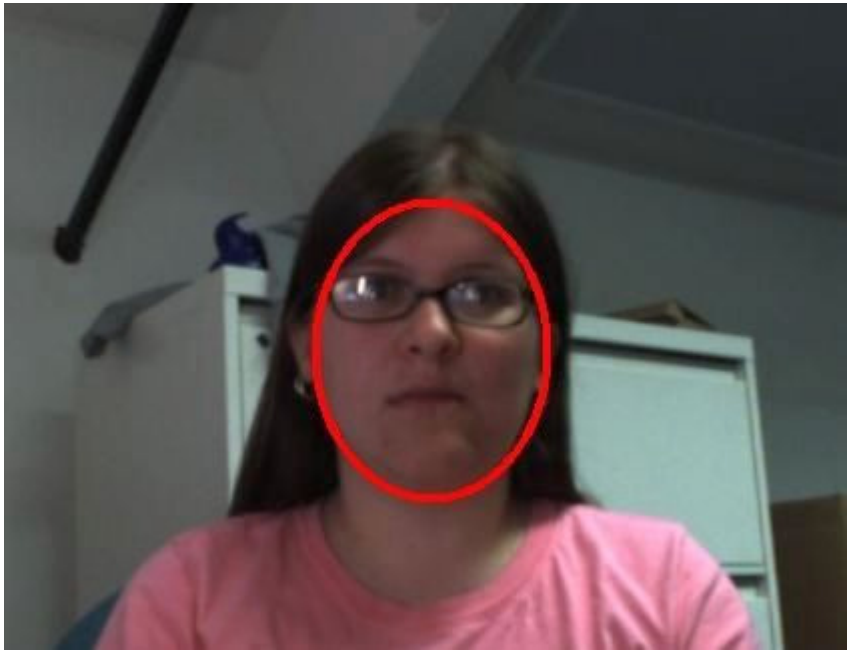

- Performed on left frame initially & every 20 frames
- Uses Haar classifiers
- Slow (~350 ms for one frame)

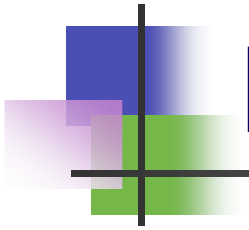

Face Detection Performance

- Requires frontal face
- Occasionally (about 5%) misidentifies
- Possible Improvements

Feature Detection

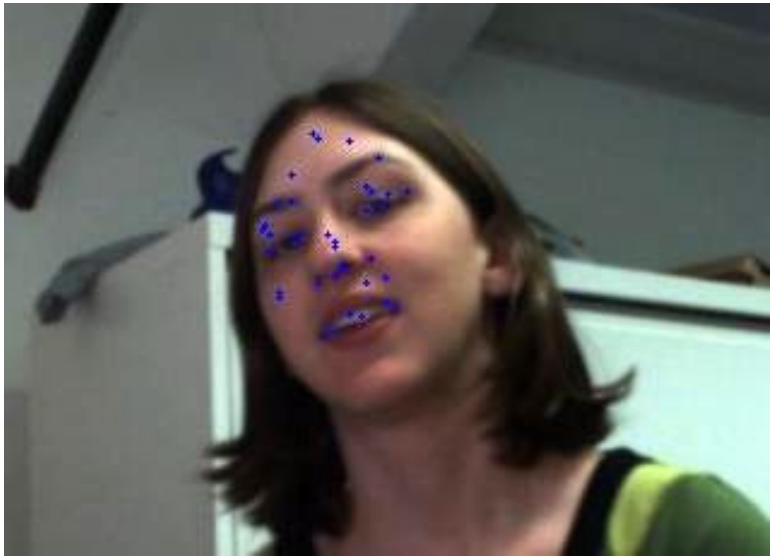

- SURF features
- Detects in a single frame
- Takes about 130 ms

Face Tracking: Original

- Mean-shift color procedure proposed by Comaniciu. et al
- Target: histogram of color distribution from initial frame
- Tracking by comparing distribution to target distribution (mean-shift)

Face Tracking: Improvements

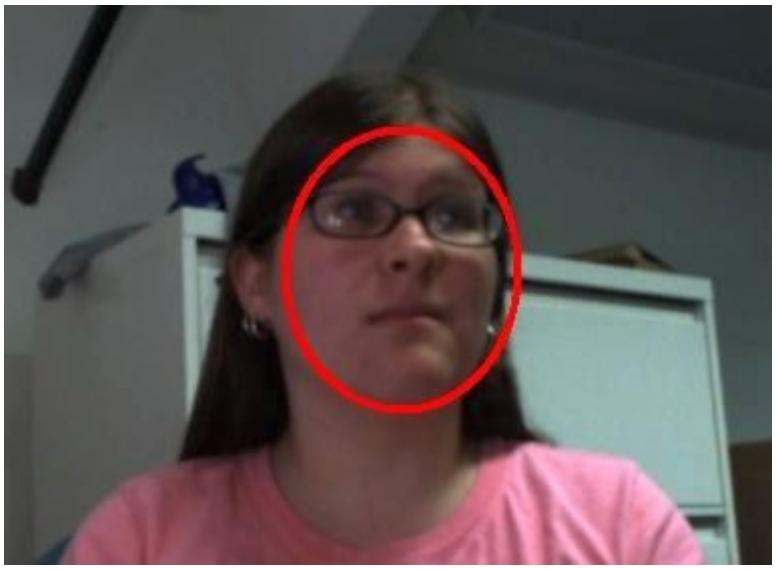

- Using two color spaces (robustness)
- Sampling (speed)
- Takes about 40 ms per frame

Face Tracking Performance

- Highly dependent on initial face
- Robust to changes in size and expression changes
- Issues with lighting changes

Feature Tracking

- Optical Flow (Lucas-Kanade)
- Performed in left frame

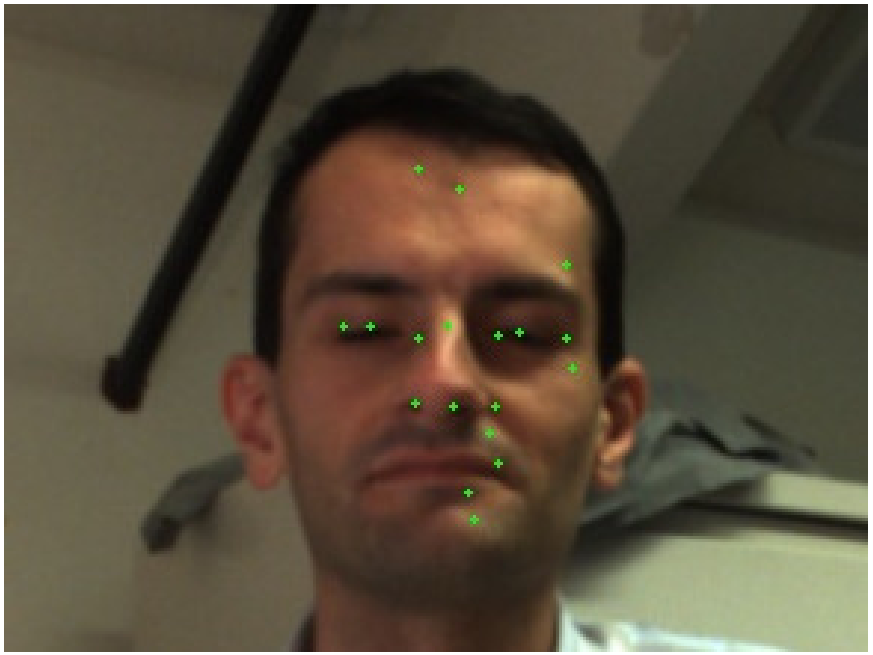

Feature Filtering

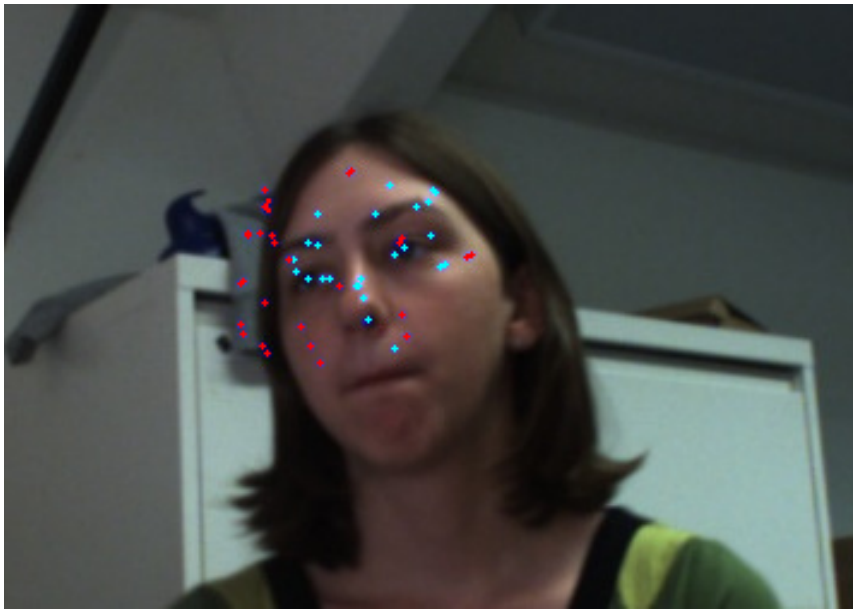

- Finds mean and standard deviation of offset (for points in face)
- Filters away points many standard deviations away from mean
- Filters away points far from face
- Signals if few points are in the face (to trigger re-detection)

Feature Correspondence

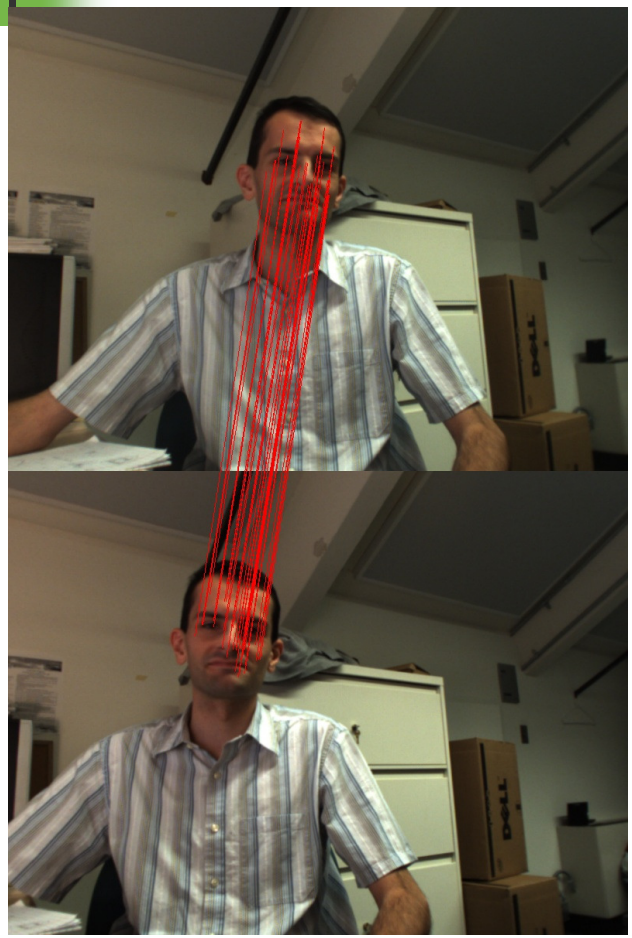

- Optical Flow (Lucas-Kanade)
- Gives matched points for pose detection
- Filters out points with high error

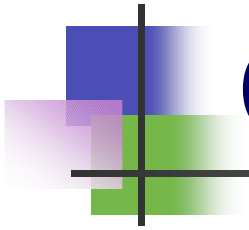

Optical Flow Performances

- Both moderately reliable without filters
- Without filters, problems with occlusions
- With filter, highly reliable

Pose Approximation

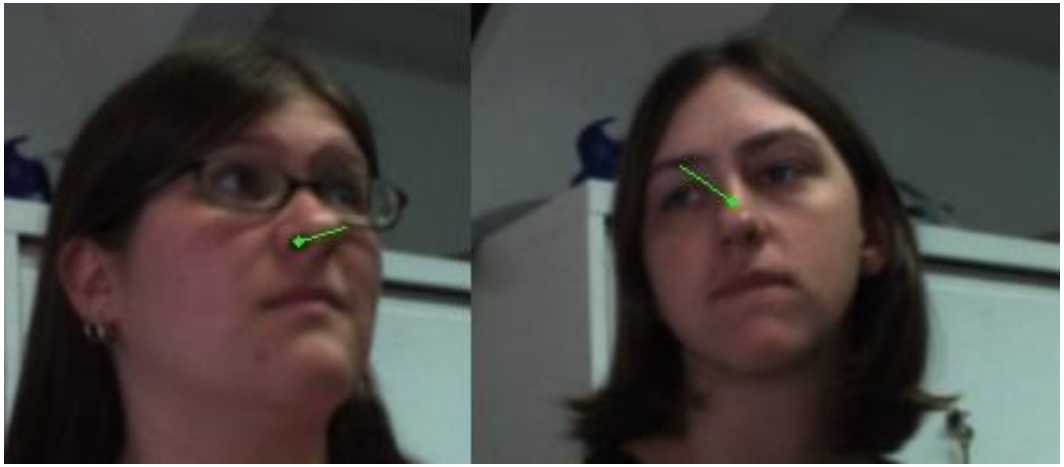

- Fits a plane to 3D points
- Normal of plane = approximate direction of face

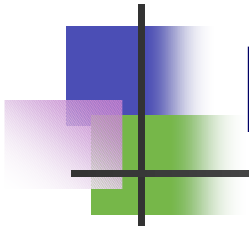

Pose Performance

- Best on well-distributed features
- Issues with poorly-distributed features
- Possible improvements

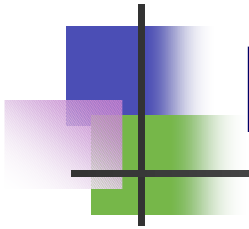

Efficiency/Speed

Tested on 800x640 image with face about 100x100.

Face Detection:	350 ms
Face Tracking:	40* ms
Feature Detection:	130 ms
Feature Tracking:	20 ms
Feature Correspondence:	25 ms
Pose Approximation:	<1 ms

*easily parallelizable

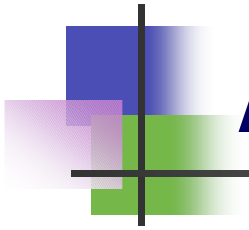

Acknowledgements

Thanks to:

Ruzena Bacjsy, Gregorij Kurillo, and everybody at the Tele-Immersive Lab for help, support, and lots of explaining.

Distributed Research Experiences for Undergraduates for funding.