




# FireFli Sensing Methodology

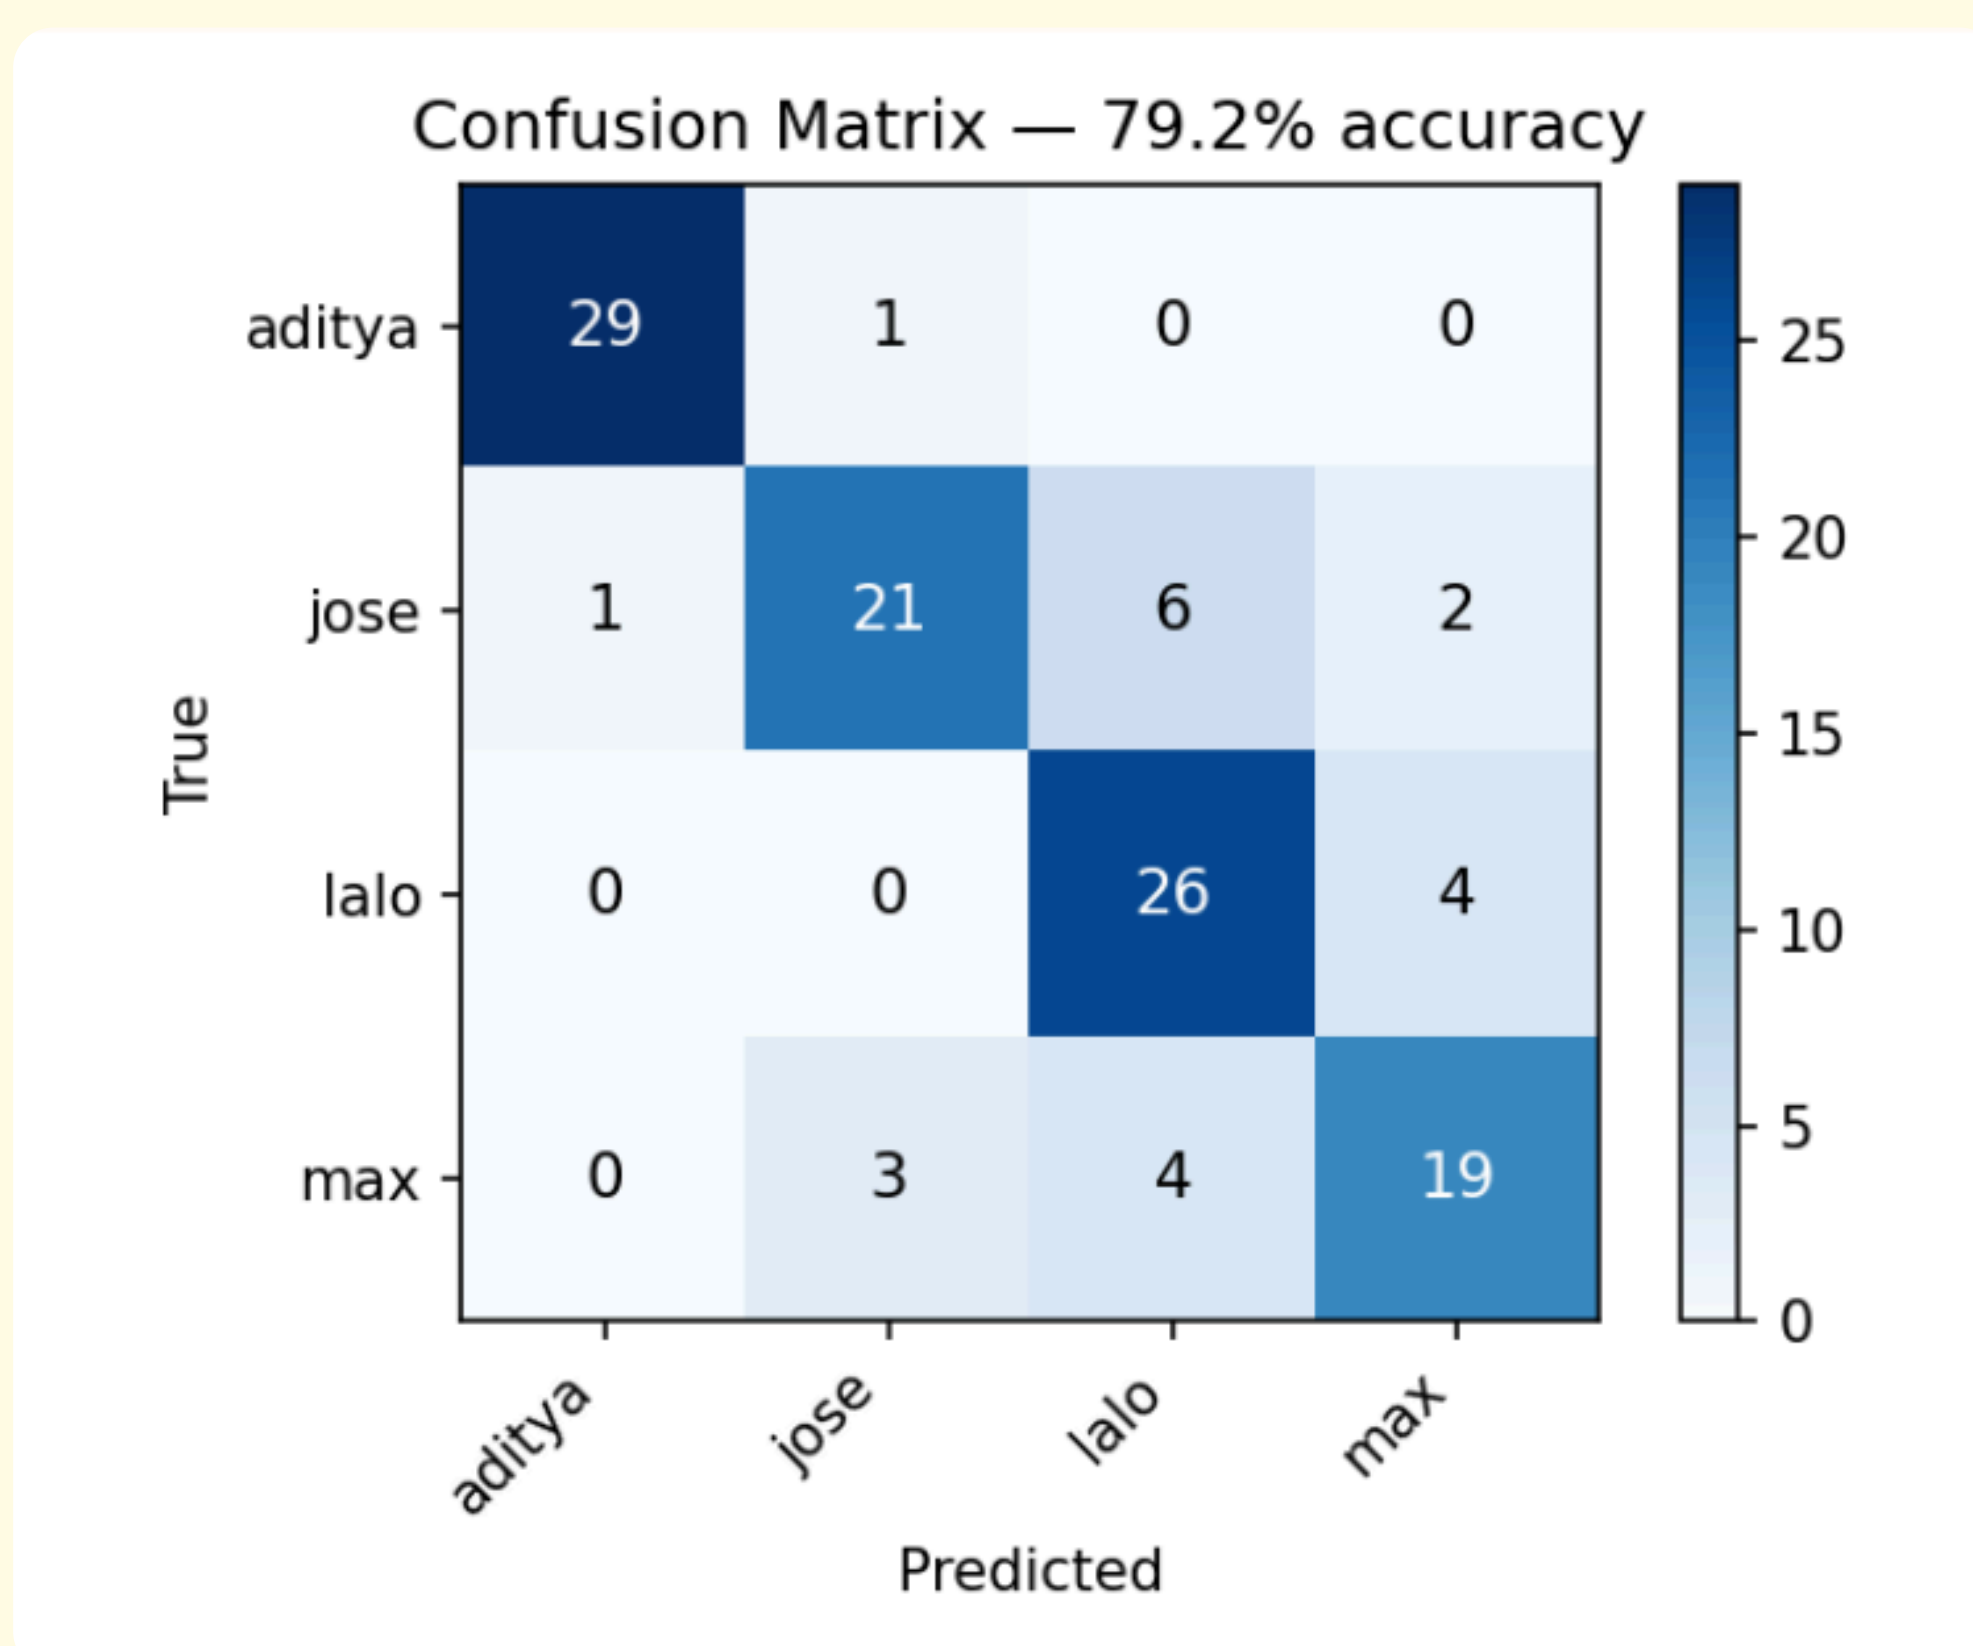
Tishya Chauhan | Kathleen May Vu | Erick Elizalde | Darius Cuenca | Ricky Cheng | Rhythm Winicour-Freeman | Mohamad Alkhabaz  
 Yuhe Shen | Joshua Hylak | Jess Bowie | Rosalind Gong | Marcellino Pena

## WiFi Subsystem

**Through Wall Detection:** Two Intel AX201 cards capture WiFi Channel State Information (CSI) via PicoScenes. A moving person perturbs the signal as Doppler shifts, which are filtered and converted to a velocity-time spectrogram via STFT. Gait features (torso velocity, stride length, cycle time, signature) are extracted from segmented walking windows and classified against known individuals using LDA.

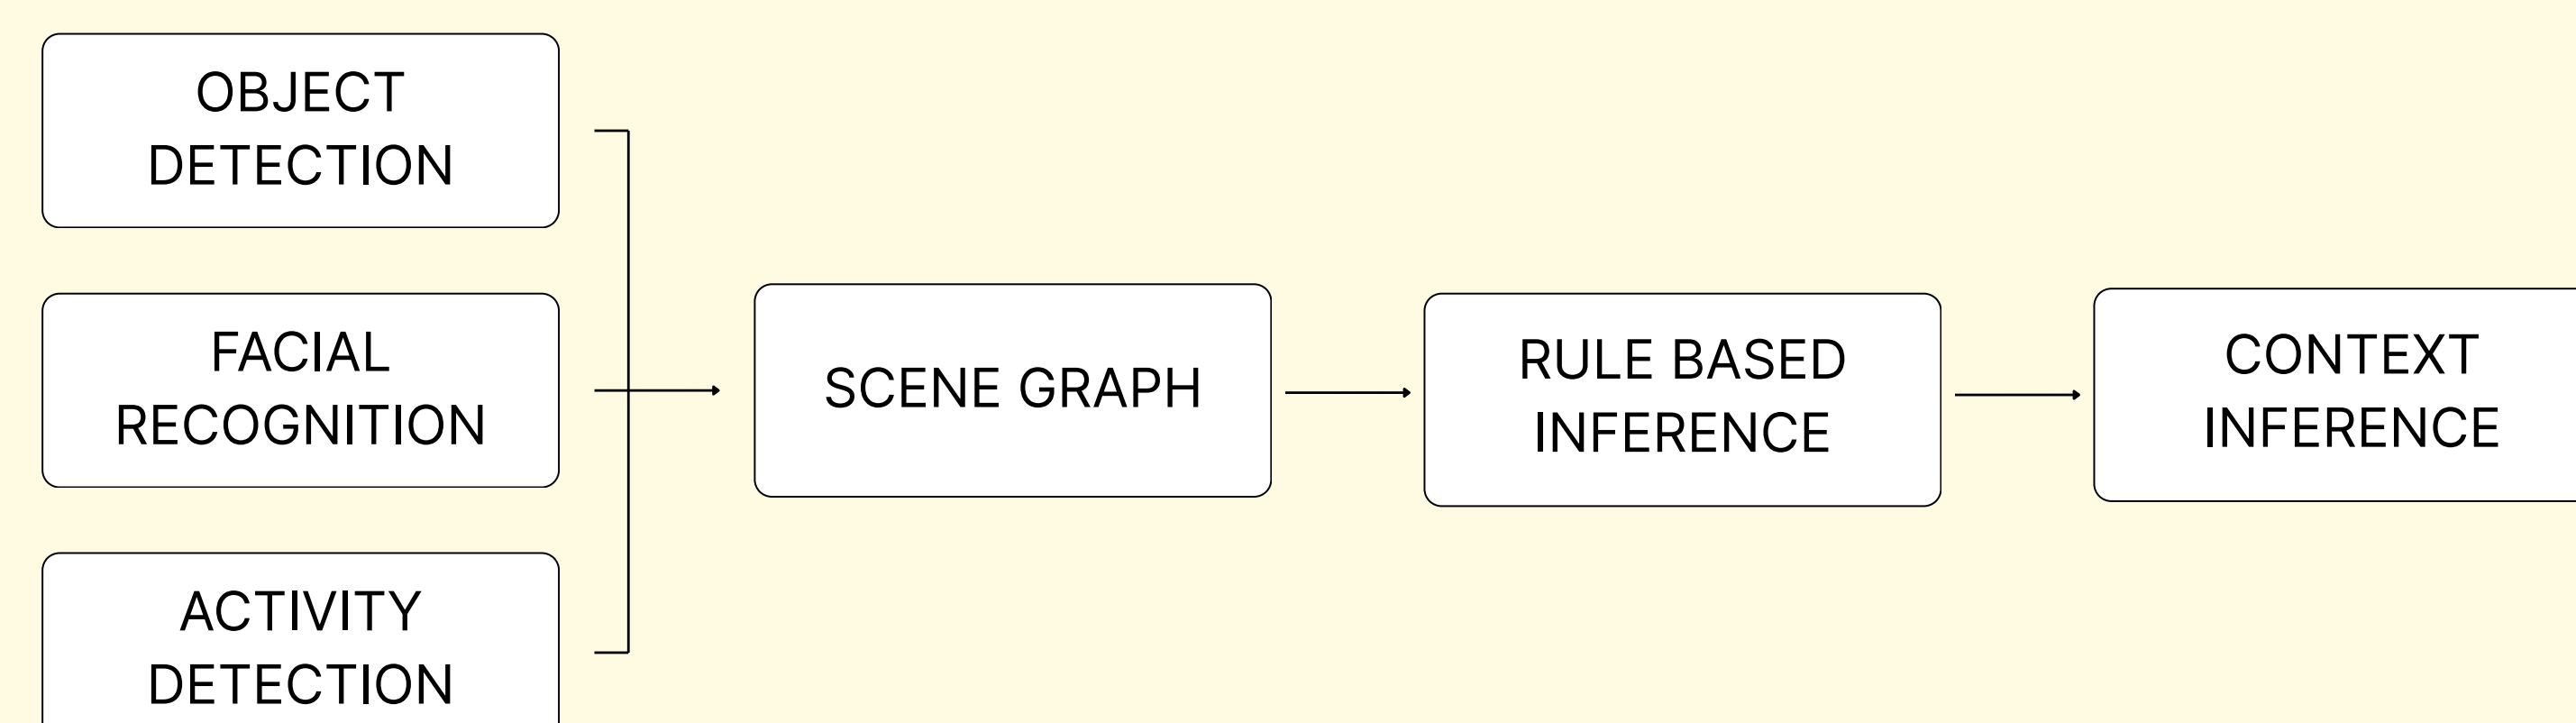


**Person Detection Accuracy:** The LDA classifier reached 79.2% accuracy across four individuals, with most errors from gait overlap between jose, lalo, and max. This confirms WiFi-derived gait enables person identification without line of sight.



## Vision Subsystem


**Context Inference:** Three detection models (object detection, facial recognition, and activity detection) run in parallel on the camera feed. Detected elements are aggregated into a scene graph, where a rule-based expert system decodes the relationships between them to produce a semantic description of the scene.



**Object Detection:** We used a YOLOv8n model pretrained on the COCO dataset consisting of 80 classes of everyday objects


**Facial Recognition:** We used InsightFace's buffalo\_l model, pretrained on the MS1MV2 dataset, to extract a discriminative embedding for each team member. Using cosine similarity, detected faces are identified.

**Activity Detection:** We used YOLO26 pose estimation to extract 17 keypoints, a Kalman filter to smooth and track keypoints across frames, and BiLSTM to classify 5 activities: running, walking, sitting, falling, and laying down.




## Navigation Subsystem

**Autonomous Navigation:** Built in ROS2 on a TurtleBot4 with onboard LiDAR, it operates in four stages: it maps the space as a 2D occupancy grid using SLAM, expands that map through frontier-based exploration (selecting frontier centroids between mapped and unknown regions), localizes its pose within the map, and plans and executes collision-free paths to target locations using the Nav2 stack with A\* global path planning, patrolling continuously and monitoring for events.



**Accuracy Trials:** To evaluate how precisely the robot reaches a target, we used coordinate transforms to locate a target object, set a goal distance of 1 m (100 cm) from it, and measured the robot's actual stopping distance. Across 20 trials, the mean stopping distance was 100.3 cm with a mean absolute error of 4.13 cm. 13 of 20 trials stopped within 5 cm of the target, and all 20 stopped within 10 cm.



## Acknowledgements:

We would like to thank Profs. Yasamin Mostofi and Ilan Ben-Yaacov, Aaditya Prakash Kattekola, and Karolina Low.