

IntelliCane: An Agentic Approach to Real-Time Obstacle Avoidance and Intelligent Decision-Making for the Visually Impaired Through a Monocular Servo-Guided Cane Using Deep Learning-Based Environmental Mapping

Chitlur, Tanay (School: Westview High School)

Ragam, Akash (School: Jesuit High School)

Over 250 million visually impaired individuals face significant challenges in navigation, limiting mobility and independence. Traditional aids like white canes and guide dogs offer minimal environmental feedback or require extensive training. As the first-ever intelligent goal-driven self-navigating aid, IntelliCane revolutionizes mobility by integrating object-aware perception with intelligent decision-making. Users can input natural-language goals—such as “go to a water fountain”—which serve as the basis for IntelliCane’s decisions. IntelliCane is composed of three core agents—Path Affordance, Object Semantics, and Depth Map Creation—supported by preprocessing tools, a centralized LLM-based aggregator, and output executors that together provide active servo guidance and verbal instructions. IntelliCane interprets visual and spatial inputs using a monocular camera system, processed through custom models optimized for embedded hardware. The Depth Map Creation Agent works in tandem with a Mapping + Pathing Tool to generate navigable paths around obstacles, while the Path Affordance Agent uses OpenCV edge detectors and camera geometry to align the user with sidewalks, trails, or hallways. The Object Semantics Agent uses transfer-trained YOLO models to detect features such as stairs, curbs, stop signs, and crosswalks, providing semantic context for goal-aware navigation. At the system’s core, a two-stage LLM Aggregator fuses agent data and user intent. The first LLM generates scene summaries, while the second uses these summaries and goals to produce real-world movement commands through direct matches or indirect reasoning. IntelliCane redefines navigation for the visually impaired, delivering intelligent autonomy that adapts to both user intent and environmental context.

Awards Won:

