Eye Exam Robotics

Intelligent Motion Laboratory, University of Illinois at Urbana-Champaign
Feb 2024 – Sep 2024

Abstract: In this study, researchers from University of Illinois at Urbana-Champaign and Duke University collaborated to develop a robot-based ophthalmic examination system. Existing robot-based ophthalmic platforms have relied on optical scanning technologies operating at relatively safe distances; however, for more precise diagnostics, it became necessary for the robot to approach the patient's eye more closely. To address this need, we aim to design and implement a high-precision motion-control scheme that allows the robot to come within 2cm of the eye, and build a system augmented with enhanced safety and reliability features.

In particular, to minimize risks associated with close proximity, we randomly introduced variables such as different skin tones, facial features, head shapes, and partial facial occlusions in conditions. By analyzing performance across these diverse scenarios, we identified inherent biases in the existing algorithm and implemented corrective measures to mitigate them.

This research was funded by the National Institutes of Health, with the primary goal of improving the system's precision and extending its reliability. We are currently implementing an updated algorithm with increased accuracy, and plan to validate overall system performance through upcoming clinical trials. Once clinically validated and deployed in real-world ophthalmic practice, this system is expected to enable faster and more widespread screening, allowing a larger number of patients to receive timely, appropriate medical care.

My Role: I used Python to conduct quantitative and qualitative stress tests on the existing face-tracking system in order to identify its performance limitations and to drive improvements aimed at increasing reliability under varied operating conditions. In addition to leveraging publicly available datasets such as AFLW-2000-3D, I collected supplemental images and performed manual annotations to create a more precise experimental environment. Through these experiments, I discovered that tracking accuracy degraded significantly whenever more than 15% of the face was occluded by the robot or other objects, or when over 40% of the face left the camera's field of view.

To address these issues, I positioned three ZED cameras at optimal locations so that at least one camera would always have a clear view of the face. I also developed a fusedtracking algorithm that combines data from all three cameras, ensuring that if one camera fails to track, another can seamlessly continue. Additionally, I implemented interpolation methods to maintain continuous tracking when all cameras miss certain facial landmarks simultaneously.

During testing, I observed notable performance degradation for certain demographic groups (e.g., specific ethnicities, children, and elderly subjects) and under different environmental conditions. To quantify these effects, I assembled a diverse dataset encompassing variations in gender, ethnicity, and age, and conducted a thorough analysis of system performance across these factors. From this analysis, I was able to define the threshold conditions at which tracking begins to fail and then updated the system design to either avoid or minimize those conditions. I also discovered that lighting changes had a significant impact on tracking accuracy, so I designed additional lighting configuration to stabilize performance under variable illumination.

Finally, I identified that the default 3D depth-estimation routine provided by the ZED camera's software development kit was not sufficiently accurate. To resolve this, I wrote a custom 3D depth-estimation module in C++ using ROS, then performed comparative tests between the ZED's built-in calibration and a four-point camera calibration method. Adopting the four-point calibration (which yielded higher accuracy) ultimately improved the system's depth-estimation accuracy by approximately 5%.

By Dahui Song.