

We first address issues dealing with distributed visual processing for a personal service robot in the Intelligent Home environment. We present an efficient and reliable framework to organize and coordinate the vision sensor nodes: fixed cameras mounted on walls, and cameras on the mobile robot. We also present key visual functionalities necessary for the robot to perform its activities. They include people detection and identification, action recognition, gesture recognition, and self-localization. We present our implementation within this framework, validated with experimental results. We then discuss infrastructure needs to manage a large number of cameras, and scalability with respect to bandwidth and processing power. Finally, we discuss the requirements for "after the fact" forensic analysis of a large corpus of video data.