

Lewis the Robotic Photographer

Cindy M. Grimm* William D. Smart†
Department of Computer Science
Washington University in St. Louis
St. Louis, MO 63105
United States

1 Introduction

Lewis is a (short) human-sized mobile robot that wanders through crowded rooms, taking pictures of people, much like a photographer at a wedding reception does. The goal is take high-quality, well-composed photographs of people non-intrusively, and to offer these pictures as keep-sakes of the conference.

2 The Exhibit

The exhibit space consists of an area for posed portrait shots and an open area where Lewis interacts with the public as they pass by the installation. Lewis will wander about this open area, looking for good pictures to take. Once a good picture opportunity is detected, Lewis will maneuver into place, take the picture and transmit it to a display workstation at the side of the exhibit. Subjects can then use this workstation to browse the pictures that Lewis has taken, choosing their favorites to be printed out as a keep-sake of their visit. In addition to simply printing out the pictures that were taken, users have the option of applying simple filters to the photographs to, for instance, have them rendered in the style of their favorite impressionist painter.

In addition to being a wandering photographer, Lewis can also take more formal portrait shots of individuals or groups. Subjects can sign up in advance for a formal sitting, where Lewis will take their portraits against a variety of scenic backdrops.

3 Robot-Human Interaction

One of the main goals of this project is to explore how humans interact with computers outside of the usual desktop setting. One of the first successful mobile robots to be deployed in a public setting was Minerva [Thrun et al. 2000], which gave pre-set tours in the Smithsonian Museum of American History in Washington, D.C. In addition to the traditional problem of navigation through a complex environment, Minerva also had to deal effectively with the general public. This general public, it turned out, was not particularly robot-friendly. For example, children tend to try climbing on the moving robot, and adults purposively block sensors or stand in the robot's path to see what it will do. Thus, dealing the a dynamic human-populated environment proved to be much more difficult than dealing with a static, mostly unpopulated one.

Lewis must also deal with the problems of interacting with humans who may not be accustomed to robots. However, the quality of this interaction is different from Minerva's, mostly because of their different purposes. As a photographer, Lewis must try to navigate through the environment and take pictures without getting in the way, or ideally, without being noticed. This is different from Minerva, whose purpose was to be a tour-guide, and thus to be the center of attention.

Lewis relies mostly on vocal requests to get a person's attention. It has been shown that using specific tones of voice (polite, mildly annoyed, *etc.*) increases the chances that humans will interact with the robot [Schulte et al. 1999]. It has also been our experience that directing the camera towards the subject as the robot utters a phrase greatly increases the quality of the interaction. We suspect that using emotive voice tones and moving the camera make the robot appear more like a living thing, allowing a human subject to more easily ascribe real intention to it. This, in turn, makes it easier for humans to interact with the robot, since they can use a human-to-human mode of interaction (which they have), rather than a human-to-intelligent-machine mode (which they likely do not have).

These techniques are used to move people if they are blocking the robot's path ("Please may I get through?"), convince people to have their picture taken ("May I take your photograph?") and to time shots ("Say cheese!"). During the installation, we plan to experiment with different specific methods of interaction, in order to determine which ones work best in which situations.

4 The Technology

Lewis uses a set of basic navigation routines to localize itself within the installation, and to safely navigate through crowds of people. Using a small number of special landmarks, the robot keeps track of its location, allowing it to go to a specific point in the installation, such as the portrait shot area, easily.

As it travels about the installation, Lewis continually takes pictures with its video camera and searches them for potential faces. Once a suitable candidate face (or faces) has been found, Lewis calculates how it has to move in order to take a well-framed shot of that person. This interfaces with the low-level navigation routines to safely move Lewis to the appropriate position.

Potential faces are chosen as candidates for a photograph based on a very simple set of composition rules that define what constitutes a "good" picture. Once we have identified such a good picture, moved into place and taken it, the image is sent to a workstation for further processing and display.

References

- SCHULTE, J., ROSENBERG, C., AND THRUN, S. 1999. Spontaneous short-term interaction with mobile robots in public places. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*.
- THRUN, S., BEETZ, M., BENNEWITZ, M., BURGARD, W., CREMERS, A. B., DELLAERT, F., FOX, D., HÄHNEL, D., ROSENBERG, C., ROY, N., SCHULTE, J., AND SCHULZ, D. 2000. Probabilistic algorithms and the interactive museum tour-guide robot Minerva. *International Journal of Robotics Research* 19, 11, 972–999.

*Email: cmg@cs.wustl.edu

†Email: wds@cs.wustl.edu