A Patient Autonomous Knowledge Sharing System For Outdoor Public Spaces

Marc Böhlen, Natalie Tan

Department of Media Study, SUNY Buffalo
marcbohlen@acm.org
natalietan@lycos.com

Abstract

This paper describes the outdoor installation "UNSEEEN", in the Reford Gardens of Eastern Québec. It emphasizes issues of autonomous self-modifying behavior during its three-month period of continuous operation. The system collected real time data on plants, insects and animals through a multi-camera vision system, interpreted this data, and shared the results with visitors. The system visualized the garden's dynamics and autonomously built hypotheses on the plant life it observed.

Introduction

Located on the Gaspé Peninsula in Grand Métis by the St. Lawrence River in Eastern Québec, the International Garden Festival of Grand Métis hosts an annual international festival of landscape and garden design. Each year, a juried group of designers, artists and architects is invited to conceive and build installations in the gardens. In 2003, UNSEEN was one of the installations selected. The exhibit began in June, ended in September, and attracted over 90,000 visitors.

In this text we describe "UNSEEN", a knowledge sharing system for outdoor spaces. In particular, we focus on questions of extended operation and robust autonomy of an electronic media art installation for outdoor public spaces.

Like a nature interpretation center, our fully automated machine-vision based outdoor installation showed visitors interesting facts about the area they were visiting. Unlike a nature interpretation center, our installation was very patient and spent hours simply watching plants.

Site Design

The design consisted of nine planting areas with ten different plants, a rest bench and an observation post (see figures 1, 2). Of these nine areas, eight were arranged in parallel strips and under the purview of a four-camera machine vision system.

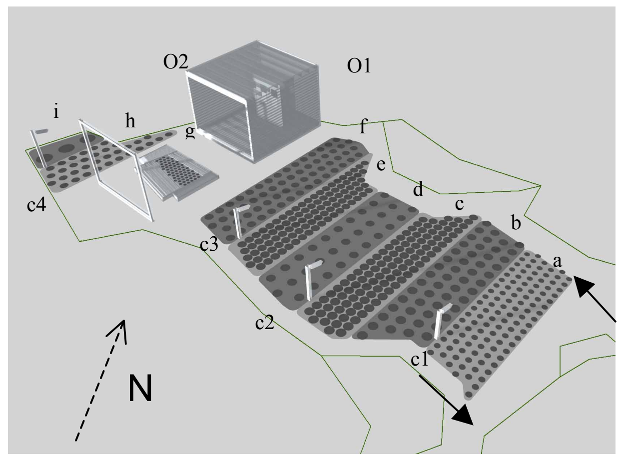

Fig. 1 CAD overview of site: a: Low bush blueberry, b: Foamflower, c: Bluebell, d: Wild sarsaparilla, e: Wild and Garden columbine, f: Canadian burnet, g: Canadian violet, h: Alpine fern, i: Dogwood; c1-c4: Cameras 1 to 4, O1: Observation post, O2: Electronics box. Arrows indicate direction of visitor traffic.

To reach the observation post, visitors walked by six of the eight strips, labeled with the proper plant names, in preparation of things to come. Rather than choose plants for visual splendor, we chose plants for biochemical, medicinal and ethno botanical properties. The sarsaparilla, for example, had been harvested by the indigenous peoples of the Gaspé and other regions and cooked into a tasty brew long before its root became the basis for root beer soda.

Our goal was to provide a novel way of experiencing a public garden and to invite visitors to reconsider their preconceptions about plants. Our methods included integrating the dynamics of the growth of plants, as measured by a multi-camera vision system, as input into a nature interpretation center with the patience to watch a garden over the course of an entire summer. But our system was not only patient, it had “second thoughts”.

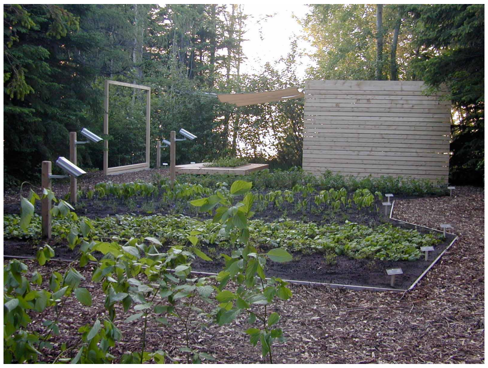

Fig. 2 The entry point of the installation site at the beginning of the festival

Three Modes of Operation

Our system operated in three distinct modes. In the first mode, all four cameras were synched and checked for the highest degree of activity amongst the plant strips. Here, “activity” was evaluated numerically as a degree of change in successive images. Motion of objects across the field of view could generate such change, for example. Whichever camera registered the highest degree of change temporarily received the complete attention of the computing system. The unaltered image stream (see figure 3) was then displayed on all three screens simultaneously in real time. Most visitors could relate to this result immediately and compare and question their own observations with this image steam. Watching things that hardly change draws one’s attention to the smallest details.
After a few minutes of streaming such data, the system rechecked its priorities, repeated this process or moved to the second mode of operation.

Fig. 3 Output from the real time image stream

A variation of this approach showed visitors the cumulative change as a gray scale image. Such an image slowly intensified as the movements under the camera accumulated (see figure 4). This mode of seeing is unique to the machine. Using differential images, one can reduce the numeric representation of change to minimal information. The method was quite sensitive and could easily register the flight path of a large bee, for example.

The second mode of observing was a comparative mode. Here we compared the current state of the garden to previous states. The system took an image from each garden section in the afternoon, when the light was soft, and stored it to disk. The collected images, four per display, were then shown in a tiled format over all three screens. This sequence represented the history of the garden in 12 episodes, visible at a glance. Over the course of the summer, the intervals between the select images increased, the difference between each image grew, but the total number of displayed images remained constant.

Fig. 4 Cumulative motion map

The third mode of observing was technically the most complex. This mode included tallying the number and density of plants in each section (as far as they were visible to the camera system). At the beginning of the festival, a given plant was the exclusive occupant of its strip. Some of the original plants achieved dense coverage of their area; others were less successful. Over time other plants moved into the territory.
In order to evaluate this process, we detected instances of each plant in each strip. The detection was based on both the color and the geometry of the plants. Color was used only to detect flowering. The colors of the blossoms ranged from red to dark blue, purple and white. The hue detection algorithm included a combination of up to three distinct hues with the option of excluding other hues. Geometric properties were gathered mainly from the leaf forms and the arrangement of leaves on the stem. The plants with large leaves, such as the Dogwood or the Sarsaparilla were robustly detected with this scheme. The plant with the smallest leaves and the smallest flowers, the low bush blueberry, was the most difficult to detect. Here

we used the geometry of the total plant and built a general model based on it[1].

The plant models at the onset of the festival were reference points with which to compare successive plant models. The models were updated daily and the differences from the original models recorded. This information, together with the number of instances of plants found and the overall density of plant material were used as weighting factors for the text generator. Thus measurable dynamics of the garden became the driving force for the way the plants were described to the visitors.

While the Δ plant model was small, the plants were described using standard Linnaean classification, followed by a factual account of measurable properties such as appearance, preferred soil type, size, seed count, etc. In this regard our system functioned as an outdoor information kiosk with factual but canned knowledge. However, the text generator was programmed to augment this factual knowledge with increasingly hypothetical statements if a given plant was found to have grown sharply in number, or to have increased or declined in density compared to the original state. Texts were built by accessing a searchable database and assembling short paragraphs on sentence level. In this state the system focused not on what was known, but on what could be imagined.

The database contained hand compiled material from published articles on all the plants, and on some animals and insects known to eat or pollinate them. We made extensive use of professional knowledge about the plants. Each plant was represented by excerpts from up to 12 articles from publications such as: *Journal of Ecology, Canadian Journal of Zoology, Communications in Soil Science and Plant Analysis, Canadian Journal of Botany, Oikos – A Journal of Ecology, The Canadian Field Naturalist, Evolution, and Systematic Entomology*. The resultant database was arranged by plant and by topic, such that the paragraphs composed from the material were valid isolated statements and in combinations were likely to make sense. In order to be able to move between pure factual information and hypothetical information, each of the sentences was weighted. The information from the plant model and density finder acted as an index for these weights. An increase in change mapped to higher indexes of 'hypothesis potential', shifting the text from the standard through the hypothetical to the imagined state.

The resulting material was a kind of tentative interpretation of many expert voices, based on real observation. This was a bona fide experimental approach, albeit one with a qualitative outcome. Visitors might find such messages confusing, or as an incentive to refer to other 'experts', or to rely on their own assessment. Either way, our system changed their experience and challenged them to take a position. Figure 5 gives an overview of the text generator. A more detailed description of the text generator and the user experiences is given in another context [Böhlen and Tan 2004].

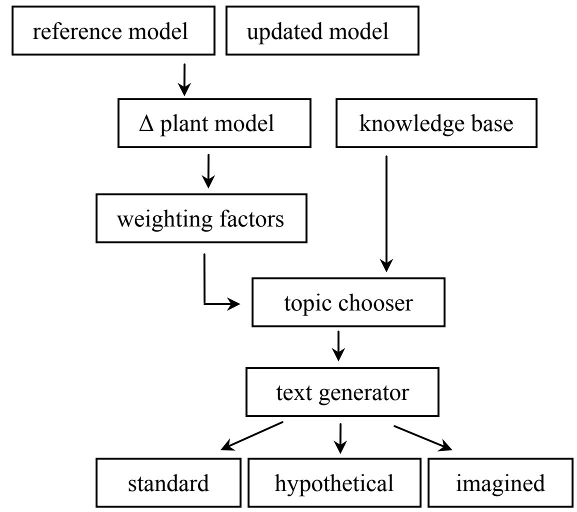

Fig. 5 Elements of the text generator

An Autonomous Vision-based System for Long-term Outdoor Activity in Public Spaces

Working with unstructured spaces creates encounters with several levels of engineering and design challenges. Outdoor public environments are not the ideal location for active electronic equipment. Since our system was placed in a large garden in the middle of a forest one-half a mile from the nearest power outlet, we had to address a variety of "adverse" conditions: line voltage fluctuations, intense temperature variations, and high levels of humidity. Since our system was designed for public access, we had to anticipate peak traffic during weekends and allay current anxieties about surveillance technologies, all the while ensuring that we could collect high quality image material.

Cameras in public spaces of leisure

The presence of high-bandwidth data collection machinery in public spaces of leisure makes some people feel uncomfortable. While most people have grown accustomed to the presence of surveillance technologies in the urban context, they expect the outdoors to be free of potential intrusion. The challenge for the authors consisted in finding ways to counter such current anxieties and still be able to collect useful data over the course of the summer. The solution was simple but effective. We placed the cameras all below (human) eye level and cast the view downwards toward the ground. This posture of downcast

[1] We integrated a commercial geometry-building library (Matrox Imaging) into our detection algorithm.

stance signaled to visitors that the vision system was watching not them, but the garden (see figure 6). This universally understood posture worked well to convince people that they were not being scrutinized.

Fig. 6 Camera position and a simulated view

Anticipating seasonal change

We modeled each camera and its position in the garden. We experimentally varied lens parameters and the path of the afternoon sun over the course of the summer until we found camera positions and lens parameters that provided good results at the beginning of the event (when the plants were short) and promised good results at the end of the event (when the plants would be at their respective full heights).

Design elements as robust landmark cues

We made use of the geometric configuration of the garden layout and paths to keep track of our areas of visual interest. The boundaries between the plant beds and the woodchip walkways were marked with sun-bleached strips of pine (see figure 7). This kind of separation is common in garden design. It worked well for the many visitors who enjoyed the soft feel of the thick woodchip carpet. It also worked well for the vision system. The garden zone was substantially darker, the wood chips created a unique texture, and the pine wood stripes stood out as a straight line under edge detection.

Fig. 7 Boundary between woodchip path and planting area

Image quality control

A major challenge for real-world object recognition is the dynamic nature of the environmental conditions with respect to illumination and visibility. The outdoor garden experienced drastic changes at different times due to varying weather conditions, and - over the course of the summer- due to seasonal fluctuations. Some researchers have addressed similar issues [Spengler and Schiele 2003] and suggested methods of recovering from errors. For this project, we implemented a simple but robust scheme of image quality control. Our goal was to have consistent image material not only for analysis but also for public display.

All four cameras were housed in enclosures that provided waterproofing and shade at all times of day. Each camera had dc-motorized active iris control. In addition to this hardware based ambient light control, we implemented a software based hue, saturation and contrast control to ensure that our image material was consistent with regards to image hue, saturation, and brightness. At the onset of each image capture sequence, we adjusted the camera thresholds mentioned above by statistically evaluating the luminescence band of a given test image and comparing the results with a reference image. This process was repeated until the images were deemed to be acceptable and ready for analysis, display or storage. Figure 8 gives an overview of this image quality control algorithm.

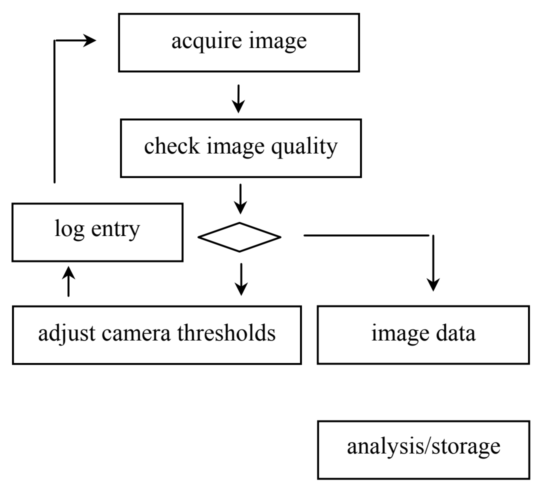

Fig. 8 Schematic of the image quality control algorithm

Poor results after rain showers. Under most lighting conditions, our system was able to find the anticipated plant forms under its area of observation. False positive finds, as shown in figure 9, occurred about 20% of the time, more frequently, however, after rain showers. While the plant coverage was sparse, the ground was visible and - when wet – was glistening and difficult to filter. Our system also performed poorly when sparse clouds moved quickly overhead and created a patchwork of high contrast light and dark patches. Luckily, these two situations were exceptions to the rule of mostly uniform lighting conditions.

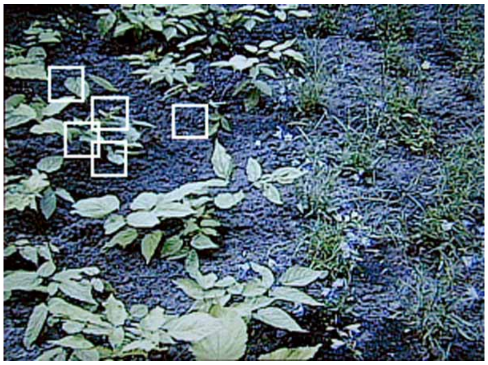

Fig. 9 Example of the system finding instances of wild sarsaparilla. The right-most white box is an erroneous result.

Anticipating hardware failure

Our narration system expected reliable input from all four cameras over the entire course of the festival. Anticipating the unlucky event of a camera failure, we added a camera functionality check to the software architecture. Prior to requesting information from any of the cameras, the system ensured that the hardware was functioning properly. In the event of a failure, the narration system was instructed to skip the particular camera and garden section, and to work with the available resources.

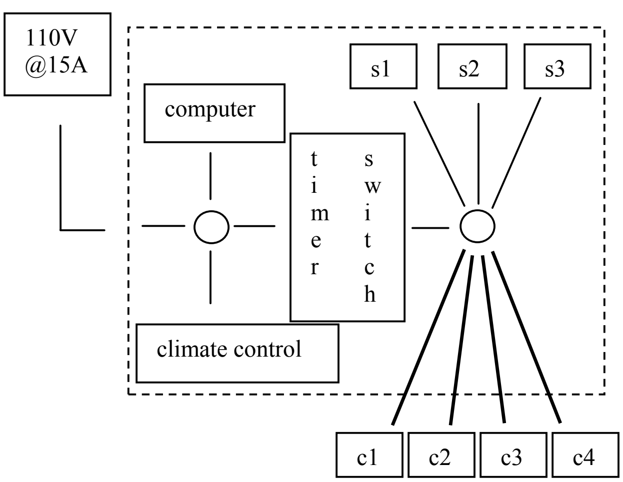

Fig. 10 Schematic of climate controlled electronics box with screens s1, s2, s3 and burial grade video cable connected to the four external cameras c1 to c4.

Simple climate control for outdoor electronics

The narration system then acted as if the problematic garden section did not exist. This short-term safety feature proved itself useful when a dc-motorized active iris on one camera failed during the first week of operation.

All electrical components were placed in a waterproof box inside of the observation post. The architecture together with a fluttering cloth canopy was designed to minimize direct sunlight from falling onto the displays during the entire day and over the entire season. Temperature and humidity were controlled by the periodic action of a strong fan that pulled in cool air from under the wooden floor of the construction. All cameras were connected to the computer's graphics capture card with outdoor burial grade video cables buried 6 inches into the ground.

An artificial custodian

In an attempt to ensure robust operation over the summer, from mid-June to mid-September, we let our system "rest" after normal visitor hours. An external timer synched with the clock of the computer was programmed to shut down the cameras and the displays once the software stopped requesting image material. This artificial custodian scheme allowed us to limit operating times to the visiting hours without intervention from service personnel. Parameter settings and select image material were saved to disk, and a system restart in the middle of the night ensured that any small unnoticed memory leaks would not accumulate to system critical failure levels. One hour before the gardens opened to the public, the system restarted itself for the day's work.

Discussion

Integrating machines into unstructured environments and public spaces is an issue common to both electronic media arts and robotics research. Both domains share the need for responsiveness and robust behavior over extended periods of time for successful operation and believability. Also, artists are exploring information processing systems as conduits for expressive media [Bouman 2003] in ways computer scientists usually do not. Where machine vision is used, environmental variability is of major concern, whether the interaction is directed to human beings [Cipolla and Pentland 1998], or to shrubs [Southall et al 1999] and [Stenz 2001]. A further common issue is the idea of appropriate action for autonomous machines in public spaces [Schulte, Rosenberg and Thrun 1999]. Which types of information gathering should machines perform in public spaces? Which roles beyond those of the expert, the toy, the pet or the companion might robotic systems be modeled as? Researchers in Hci [Weal et al

2003], in social robotics [Fong, Nourbakhsh, and Dautenhahn 2003] and in media arts are now formulating approaches to such challenges.

We thank the patrons of the Reford Gardens of Grand Métis, Québec, for the opportunity to contribute a pleasurable and practical case study to this new field of inquiry.

References

[Böhlen and Tan 2004] M. Böhlen, N. Tan, “Garden variety pervasive computing”, IEEE Pervasive Computing, Special Edition on Pervasive Computing for Art, Design and Entertainment, Jan-Mar issue, 2004.

[Bouman 2003] M. Bouman, “The machine speaks the world’s thoughts”, Corps automates_automata, Parachute 112, pp. 109 – 125, 2003

[Cipolla and Pentland 1998] R. Cipolla and A. Pentland, “Computer vision for human-machine interaction”, Cambridge University Press, 1998.

[Fong, Nourbakhsh, and Dautenhahn 2003] T. Fong, I. Nourbakhsh, K. Dautenhahn, “A survey of socially interactive robots”, Robotics and Autonomous Systems 42, pp. 143–166, 2003.

[Schulte, Rosenberg and Thrun 1999] J. Schulte, C. Rosenberg, and S. Thrun. “Spontaneous short-term interaction with mobile robots in public places”, In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 1999.

[Spengler and Schiele 2003] M. Spengler, B. Schiele, “Towards robust multi-cue integration for visual tracking” Perceptual Computing and Computer Vision Group, ETH Zurich, Switzerland, Machine Vision and Applications, 14, pp. 50–58, 2003

[Southall et al 1999]. B., Southall, T. Hague, J. Marchant, B. Buxton, "Vision-aided outdoor navigation of an autonomous horticultural vehicle," Proc. of the First International Conference on Vision Systems (ICVS), 1999.

[Stenz 2001] A. Stentz "Robotic technologies for outdoor industrial vehicles", Proceedings of SPIE AeroSense, Unmanned Ground Vehicle Technologies, April 2001.

[Weal et al 2003] M. Weal, D. Michaelides, T. Danius, M. Thompson, K. Mark, and D. De Roure. 'The ambient wood journals - replaying the experience'. Proceedings of ACM Hypertext'03, The fourteenth conference on Hypertext and Hypermedia 2003, Nottingham, UK, 2003.