

Leveraging Limited Autonomous Mobility to Frame Attractive Group Photos

Jason Campbell^{1,2}

¹Intel Research Pittsburgh
Pittsburgh, PA USA

{jason.d.campbell,padmanabhan.s.pillai}@intel.com

Padmanabhan Pillai¹

²The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA USA

Abstract - Robot photographers have appeared in a variety of novelty settings over the past few years and typically have exploited rudimentary image-content-based approaches to identifying potential photographic subjects. These approaches are primarily limited to human subjects and further progress along content-based lines is hamstrung by slow progress on the general computer vision problem.

In this paper, we present a mobile robot system which solves the group-picture-framing problem without requiring content-based methods. The system finds photographic subjects based on measurements of motion parallax obtained via optical flow during robot movements. Our method requires only sufficient contrast to permit reasonably accurate sparse optical flow field estimation and is completely independent of any content-based image heuristics. The result is a working mobile robot system that can correctly photograph human and non-human subjects in a variety of posed-subject situations, and produce well-framed, cropped images for printing on standard-sized photo paper.

Index Terms - photography; mobile robot; computer vision; picture framing

I. BACKGROUND ON GROUP PHOTOGRAPHY AND THE PICTURE-FRAMING PROBLEM DEFINED; RELATED WORK

Group photography presents a somewhat restricted form of the photo composition problem in which the subject(s) are posed in close proximity to one another and the background is of little photographic interest. Group photo subjects generally collaborate with the photographer to produce a mutually-acceptable image, including cooperation on the physical placement of subjects (standing close together), their behavior (standing still), and displaying appropriate facial expressions (“say cheese!”).

In this context we define the “picture-framing problem” as the selection of a boundary which encloses all visible portions of the subjects and as little of the background as necessary to avoid visually crowding the subject. From an aesthetic perspective it is important that the subject(s) be centered within this boundary – that is, that each of the four boundary edges be a similar standoff-distance from the subject (to the extent permitted by the total camera field-of-view). It is also necessary to select an appropriate aspect ratio to exploit as much of the printed-image area as possible without clipping important parts of the subject. See Fig. 1 for an example of picture framing. This paper describes an algorithm which solves the picture-framing problem in common real-world environments and exploits that solution to produce visually appealing group photos.

In contrast to group photography, the composition and timing of high quality editorial, candid, and art photos depend upon fine shades of image interpretation – subtleties beyond the present state of the art in computer vision. In the present work we avoid these sorts of image-content-based decisions in order to focus our efforts on what can be achieved without image interpretation. An important consequence of this is that the robot photographer presented here performs equally well for inanimate and non-human subjects as well as for humans. Thus, this is the first robot photographer which can capture all-robot and mixed human/robot group photos just as effectively as all-human ones! Of course, as the state of the art in automated image understanding advances, hybrid image composition strategies will become possible in which the present algorithm may play only a partial role.

Previous “robot photographer” work described in [1,2,3] applied face and skin detection techniques and rudimentary principles of image composition to photograph people non-interactively in social settings. However the results depended upon subsequent human culling of “interesting” from “uninteresting” shots and might better be termed “serendipitous” than “candid”.

The Robot Photographer presented in this paper tackles a different photographic problem, demonstrating the feasibility of using depth from motion information alone to produce well-framed group photographs in posed, cooperative settings. Further contributions include a fast, simple rectangle-consensus algorithm to determine subject area from such motion data, and a set of heuristics to produce aestheti-



Fig. 1: The Robot Photographer, and an Example of Picture-Framing
The cropped, resized image at the upper right is automatically produced by the algorithm described in this paper.
The original source image is shown at the lower right

cally pleasing cropped images suitable for printing. Finally, results from an implementation of these techniques are presented and a discussion of broader applicability of the techniques is included.

II. PHYSICAL CONFIGURATION AND MODUS OPERANDI

The Robot Photographer (Fig. 1) is designed to be deployed in social situations where the availability of group photos provides participants the opportunity to create a memento of their attendance and the others who they meet at the occasion. For instance, we first deployed the Robot Photographer at our annual laboratory open house in August 2004.

A. Example Scenario

A small group gathers in front of the robot. With a mouse-click on an appropriate GUI, which appears on a display screen adjacent to the robot, the group requests a photo be taken. Using vision techniques [11], the robot ascertains that it has adequate space to maneuver and then translates a small distance to the left and to the right while observing a broad field of view in the general direction of the subjects. By identifying motion parallax in its field of view and clustering parallax velocities the robot identifies the boundaries and position of the photographic subject(s). The robot then asks the subjects to assume appropriate facial expressions (“Say Cheese!”), and shoots a full-resolution photo. Finally the robot presents the final photo to the subjects on an adjacent display screen and asks for approval to print the image.

B. Hardware Components

We use a Personal Exploration Rover [4], as shown in Fig. 1, as the mobile robot platform for the Robot Photographer. The PER robot is configured for omnidirectional movement at approximately 3cm/s. Attached to the robot’s upper surface is a high quality digital camera (Canon S60) mounted on a tripod. For purposes of hazard detection a small web-cam (Logitech QuickCam Pro IV) with a wide field of view lens (Unibrain US2042) is attached to the robot’s pan-tilt head. Vision processing for both cameras runs on a single adjacent desktop (3.6 GHz Intel® Xeon™) or laptop (1.5 GHz Intel® Pentium® M) PC, depending upon the deployment situation. Both cameras are connected via USB tether to the PC. The robot is controlled by commands transmitted over a peer-to-peer 802.11b wireless LAN connection. Although the robot carries a powerful embedded microprocessor (400 MHz Intel® PXA250), software compatibility issues with the connection to the digital camera as well as a desire to incorporate low-latency video in an off-robot user interface led us to adopt this tethered approach.

The webcam is used purely for hazard detection via visual odometry [5,11] and operates at a resolution of 640x480 pixels, achieving approximately 10 frames per second peak frame rate. All photographic subject selection and photography is done with the much higher quality digital camera, which uses a Canon-specific extension to the standard picture transfer protocol (PTP) that permits the shooting and immediate transfer of photos over the USB interface, rather than storing them into the camera’s flash memory for later retrieval. This communication is enabled through the libptp2 library [6] and custom software based

on the ptpcanon package [7]. The high-resolution camera operates either as a 2592x1944 pixel still camera or in “viewfinder” mode as a video camera achieving 5 frames per second at a 320x240 pixels. At both resolutions, JPEG compression is applied, with a high compression ratio on the viewfinder images and a lower compression ratio on the still images. Because subject selection, framing, and final high-resolution imaging are performed using the single main camera, problems with static parallax such as would arise between multiple cameras, a laser ranger and a camera, or between a high-resolution camera and a stereo rig are avoided.

Attached to the PC is a color inkjet printer (HP 5550) loaded with 5x7 inch photo paper stock. This stock has been preprinted on the reverse side with information about the Robot Photographer and the event at which the photos are being taken.

III. SUBJECT-IDENTIFICATION ALGORITHM

In order to identify photographic subjects the Robot Photographer uses a multiple-step process that depends on both visual analysis and some degree of subject cooperation (specifically, that the subjects hold reasonably still). The algorithm assumes that photographic subjects stand close together in a group, are arranged roughly parallel to the image plane of the camera, and allow some depth between themselves and unwanted background objects.

The subject identification process first produces sparse relative depth estimates for the entire field of view based on motion parallax. This is accomplished by translating the robot and attached camera laterally, i.e., along the image plane, and computing the motion parallax from multiple frames acquired while moving. Given the canonical posing of group photo subjects described above, and a perspective camera, regions corresponding to the subject will have similar parallax shifts. By clustering points based on this estimated depth, the subject regions are identified. Finally, the algorithm determines the subject crop rectangle, a best-fit bounding box that fully encloses the subject regions. The details of the subject identification process are as follows.

A. Motion Parallax Estimation

Motion parallax provides a strong cue to scene depth. Given a moving camera and static environment, parallax can be directly measured using optical flow techniques on the video stream. This is the approach taken by the Robot Photographer, which uses the Lucas-Kanade method [8], implemented in image pyramids as part of the OpenCV computer vision library [9, 10] to estimate optical flow as seen through the camera viewfinder. Even though subpixel tracking is performed, optical flow field velocities exhibit increasing noise for smaller displacements. To increase the resolution of our depth measurements we consider the cumulative optical flow displacement over five frames rather than simply the displacement between the two most recent frames. While this is in principle equivalent to choosing a wider stereo baseline or to decimating the incoming video sequence (i.e., estimating the flow field at a lower frame rate), the intermediate frames allow tracking errors, especially those due to spatial aliasing in scene areas with regular patterns, to be minimized.

For a static scene, depth is inversely proportional to relative optical flow field velocity. Regions containing points exhibiting large optical flow vectors will be closer to the camera than those with short vectors. Fig. 2 shows an example of this. The Robot Photographer does not make the inverse transformation to depth but instead directly clusters the flow velocity values, as will be described in the next section.

Limited dynamic behavior in the scene (i.e., other moving entities) can also be accommodated without confusion by observing the optical flow field over time. As particular image patches are tracked the direction of the corresponding optical flow vectors is observed over several frames. Highly inconsistent directions for the same image patch serve as an indication of a tracking error or an independently moving portion of the scene. The corresponding optical flow field elements are ignored in the clustering and rectangle calculation steps below. This step assumes a smooth motion model for the robot, but an equivalent screening step can be performed for non-smooth robot motion.

B. Subject Identification via Depth Clustering

Given the parallax displacements measured in the previous step the Robot Photographer proceeds to cluster the sparse tracking points used by the optical flow estimation process into groups of similar depth points. The parallax displacements are inversely proportional to depth, and so the clustering process can be performed directly on the displacement values with no loss of information. The clustering is performed according to the procedure in Fig. 3.

Once every displacement value has been assigned to a cluster, the clusters are sorted by decreasing displacement values (which corresponds to sorting by increasing depth). Given a canonical group photo pose, one would expect the subjects to be in the foreground (low depth values), of similar depth (within same cluster), and occupy a significant portion of the overall scene. Hence, the first cluster in the sorted set that includes more than 5% of the total number of displacement values (track points) is selected as the subject. The 5% criteria helps increase robustness of the subject selection to tracking errors and noise which may not be screened out by the flow-direction-consistency criteria discussed in the previous section.

C. Crop-Rectangle Selection

Having established a cluster of points corresponding to the photographic subject, a crop-rectangle is selected which exactly encloses all of those points.

D. Rectangle-Consensus

Because steps A-C above may execute as many as fifty times during the robot's short subject-selection movement, a mechanism is needed by which the resulting crop rectangles can be combined into a "consensus" rectangle. Using such a consensus operation further increases the robustness of the subject selection system since although individual candidate rectangles may be poorly placed due to jerky robot motion, tracking errors, subject, or background movement, the consensus value is still likely to reflect a stable subject boundary in the field of view. Common operators such as average or median would not be appropriate when used on rectangle




Fig. 2: Optical flow vector lengths vary by depth

Each circle's intensity indicates the magnitude of the motion parallax estimated at that tracking point. Lighter colors denote greater parallax whereas darker circles denote lesser motion parallax. The two track points marked with crosses (on the blank rear wall) are ignored because they have exhibited "unsmooth" flow over the recent 7 frames.

bounding coordinates, nor would integrating rectangle enclosures because the subject "moves" (due to the robot's motion) in the field of view during this process.

Instead, the Robot Photographer uses a voting scheme where crop rectangles of similar dimensions are grouped and contribute votes to an assertion that the correct subject rectangle is of that size. (See Fig. 4) After the robot's motion is complete and before the final photo is taken, the most recent rectangle from the winning group is taken to be the consensus subject rectangle.

Because the particular chosen rectangle may have been found before the end of the robot's movement, an adjustment must be applied to shift that rectangle a depth-dependent number of pixels corresponding to the robot's further movement. (See Fig. 4) This adjustment is calculated based on the remaining number of video frames during which the robot moved multiplied by the average shift per video frame observed in the winning group. Using this method of adjustment avoids the need for any explicit mapping between parallax and actual depth in the scene.

IV. IMAGE FRAMING ALGORITHM

After identifying the subject region, the robot stops moving, asks the subjects to assume appropriate facial expressions by playing a sampled sound ("Say Cheese"), and commands the digital camera to take a high resolution still photograph. Once the camera has captured and transferred

```

for each displacement value {
  look for a matching displacement cluster
    (within 30%);
  if a matching cluster is found,
    add this displacement to it,
  otherwise create a new cluster
    at this displacement value
}

```

Fig. 3: Procedure for clustering displacement values

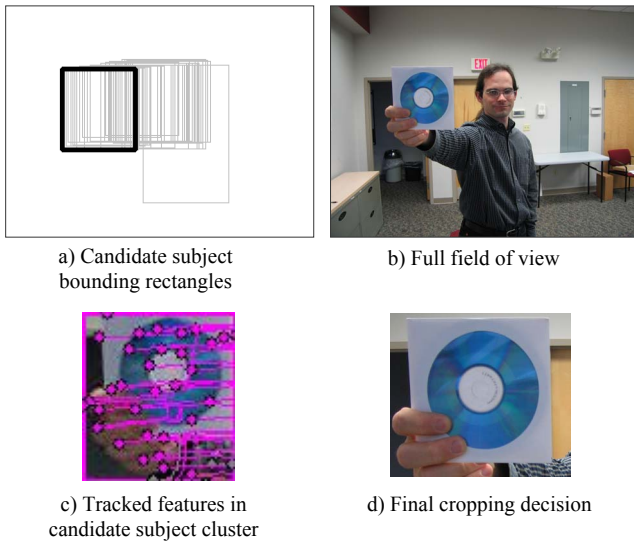


Fig. 4: Subject rectangle consensus process

Each new video frame yields an estimated bounding rectangle for the subject (based on the closest major cluster of parallax measurements). Around 50 such rectangles are estimated per subject-localization cycle. The full set of candidate rectangles for one cycle is shown in (a) and the corresponding camera field of view in (b). The candidate rectangles in (a) are spread across the frame because the subject shifts in view due to motion parallax. Because our rectangle-consensus method looks only at width and height it is unaffected by this shift. The dark black rectangle in (a) is the resulting consensus subject bounding rectangle. Inset (c) depicts the set of features being tracked on the candidate subject at the end of the cycle, and (d) shows the final cropping decision.

this image, the system produces a properly framed and cropped photograph for printing. This is a nontrivial step as the system must crop the image in an aesthetically pleasing manner that both covers the subject area, without crowding the subjects, and ensures that the image matches the output medium aspect ratio. The image framing algorithm starts with the subject crop rectangle, and modifies it to produce a properly framed cropping rectangle for printing.

A. Adding Breathing Room

Closely cropped photographs are often aesthetically unpleasing, and most human photographers will include a small amount of “breathing room” around the subjects in a photo. The Robot Photographer accomplishes this by introducing a factor to control the tightness of the crop. The crop rectangle selected in the previous step is grown by an amount corresponding to 10% of the selected rectangle’s width. This expansion is performed equally on all four sides of the crop rectangle, that is, the top and bottom of the rectangle are also expanded by the fraction of the rectangle’s *width*, not height. This ensures that the added “breathing room” is balanced at image corners. Where the expanded rectangle would exceed the camera field of view, it is truncated to the field of view rectangle. This occurs most frequently at the bottom edge of an image because many subjects extend down to the floor, below the camera’s field of view.

B. Matching Aspect-Ratios

Printed photos typically have a standard size related to the sizes of photo paper or frames available. The Robot Photographer is designed to print 4x6 inch landscape images on standard 5x7 inch photo paper with appropriate

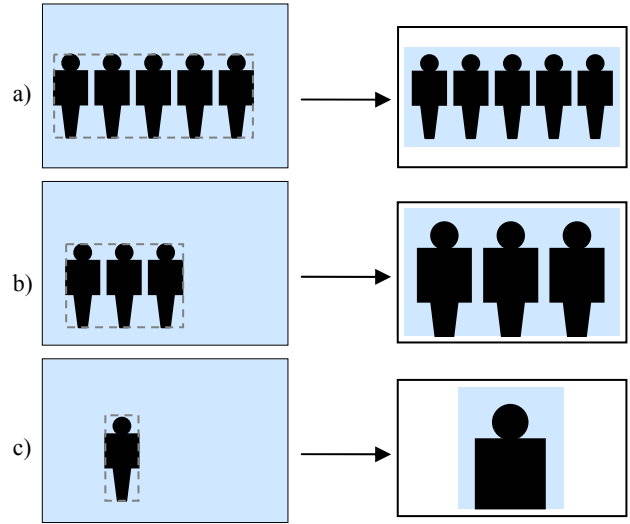


Fig. 5: Adjusting subject crop rectangle to match target aspect ratio
On the left, the filled/solid rectangles illustrate the full extent of the camera image. The dotted rectangles show the enclosing bounds around the subject established via motion parallax and optical flow.
On the right, corresponding photo print mockups illustrate the resulting cropped and scaled photo using the image framing algorithm described in Section IV.

margins and room for a caption. Ideally, the printed photo will use as much of the image area as possible, given the selected print margins and layout, with the subject reproduced as large as possible. Unfortunately, as there is no guarantee that the subject will closely match the proportions of the output medium, an adjustment of the aspect ratio (width/height ratio) of the crop rectangle may be necessary.

Beyond the adjustment for “breathing room,” the crop rectangle is never enlarged. This ensures that the subject fills the image frame and no additional distracting background clutter is added. Any adjustment to the aspect ratio is made by further cropping at the bottom edge of the subject rectangle. There are two reasons for this: First, because most subjects stand on the ground or are supported by something/someone standing on the ground, the bottom edge of the camera field of view will often be cut at an arbitrary part of the subject or its supporting apparatus. Thus, the bottom edge of the selected crop rectangle will be relatively arbitrary and hence more amenable to adjustment than the other edges. Second, in a group photograph, the subject width is not generally amenable to adjustment without cutting individuals off. Similarly, the top edge is often corresponds to the subjects’ heads and should not be cropped further. It is typically least objectionable to crop the lower extremities in images of standing subjects, especially human ones.

If the aspect ratio of the selected crop rectangle is greater than the aspect ratio of the target printed image (i.e., a wide, short subject area), then the full estimated subject rectangle is maintained, and the image is scaled to the width of the output medium. This results in additional margin above and below the image (Fig. 5a). The subject is reproduced as large as possible without distortion and without introducing foreground / background clutter.

If the aspect ratio of the crop rectangle is somewhat less than that of the target printed image (i.e., a compact subject

similar in proportions to the target photo paper), the bottom edge of the subject rectangle is further cropped to until the ratios match (Fig. 5b). This discards the lowermost portion of the subject, but ensures that the remaining subject area maximally fills the output image frame.

Finally, if the aspect ratio is substantially less than the target print ratio (i.e., the subject area is very tall and narrow), then the preceding rules may cut off too much of the subject (especially for single human subjects). Thus, if the aspect ratio of the crop rectangle is less than 2/3rds of that of the paper, the bottom edge is moved just enough to ensure that the output image aspect ratio is no smaller than the inverse of the golden mean (Fig. 5c). The image is scaled to match the height of the output medium, which results in additional margin on the sides. This last heuristic permits a tall, thin subject to be cropped reasonably and reproduced without switching to “portrait” orientation on the paper. In general, such a switch may be preferable to the above rule, but the photo paper used with the Robot Photographer is preprinted on the reverse side with information about the project and we wanted to maintain a consistent orientation on both sides of the final print.

V. COPING WITH HAZARDOUS WORKING CONDITIONS – NOT FALLING OFF THE TABLE

The Robot Photographer’s current incarnation operates on top of a table so as to better interact with humans and to photograph human subjects from a more pleasing angle. But to identify a photographic subject, the robot must move around the table. This means the robot is also in danger of falling off the table. A comparable floor-standing robot might need to instead worry about running into obstacles such as people, furniture, or walls. For both of these situations, hazard detection can be performed using motion parallax information. For our tabletop Robot Photographer we take advantage of discontinuities in the motion parallax field to detect the table edges and adjust the robot’s position accordingly before beginning a major movement. This optical-flow-based mechanism is robust and allows the robot to safely detect table edges from a variety of angles and distances.

To check for a table edge, we turn the pan/tilt webcam attached to the robot in the direction of travel, and then command the robot to move a very short distance. The vision system estimates optical flow using the images from the webcam and looks for a substantial discontinuity in flow field velocity over a short angular interval in the field of view (see Fig. 6). The consistent presence of such a discontinuity over a few frames of video is taken as a strong indication of a precipice, and within approximately 1cm of travel, the vision system can accurately determine the location of any table edge in the robot’s field of view. If the edge is found too close to the robot to permit the 15cm translational movement used in identifying photographic subjects, then the robot first maneuvers to get more space before attempting to frame the subject.

VI. OBSERVED PERFORMANCE

The Robot Photographer has been quite effective at producing well-framed and aesthetically pleasing photos of posed individuals and groups. When our assumptions of




Fig. 6: Precipice detection using optical flow vectors

Each circle’s intensity indicates the velocity of the optical flow field at that tracking point. Lighter shades indicate higher velocities and darker shades indicate lower velocities. Depth from the camera is inversely correlated to optical flow field velocity in this case and precipice detection is accomplished by looking for significant discontinuities in the flow velocities. (such as occur at the table edge in this example)

posed, still subjects with a relatively static background are valid, the Robot Photographer produces excellent results. For instance, in a controlled experiment involving 20 photograph framing cycles, including 10 attempts at photographing a single human subject and 10 attempts at photographing a pair of human subjects, the Robot Photographer correctly framed 19 of 20 of the photographs. The remaining photograph was cropped too widely by approximately 50%.

When our assumptions do not hold, photographic performance suffers. For instance, in actual use at public events the Robot Photographer can be confused by independent motion in the background, especially when that motion is smooth.

For examples where correct framing is achieved, see Figs. 7a, 7b, and 7c, as well as Fig. 1. The images in Figs. 1 and 7a are correctly framed, with the subject area slightly cropped to match the 3:2 aspect ratio of the output. Fig. 7a in particular illustrates that depth clustering and subject selection are robust to slight variations in subject distance, and correctly include both the person kneeling in the foreground and those standing behind. The image in Fig. 7b shows good framing for a tall, narrow subject, with preference shown to subject coverage and background elimination rather than matching the target aspect ratio. In Fig. 7c, the Robot Photographer has correctly distinguished the foreground subject from the individual acting as a support.

For examples of incorrect framing, see Figs. 7d, 7e, and 7f. Fig. 7d shows a situation where background movement fools the depth estimation process, and results in the background regions being identified as subject area. The problem arises because the two moving people in the background cause sufficiently large and smooth optical flow fields from their own motion that they confuse the subject identification process. The motion parallax due to independent motion is difficult to distinguish from the motion parallax due to the robot’s own motion. In the image shown in Fig. 7e, the camera field of view includes non-subjects

that are coplanar with the true subject. Due to its similar depth, the non-subject material is clustered with the subject, resulting in a poorly cropped photo. A modification to our clustering technique to account for spatial distance as well as parallax field magnitude might improve this particular case. Finally, for particularly flat subject regions, or overly cluttered backgrounds, there may be an insufficient number of tracking points on the main subject or parts of the subject. This results in either a completely misidentified subject area, or a selected subject region that only partially covers the main subject, producing over-cropped results as in Fig. 7f.

VII. CONCLUSION

This paper has introduced a mobile robot system that solves the picture framing problem by exploiting specific

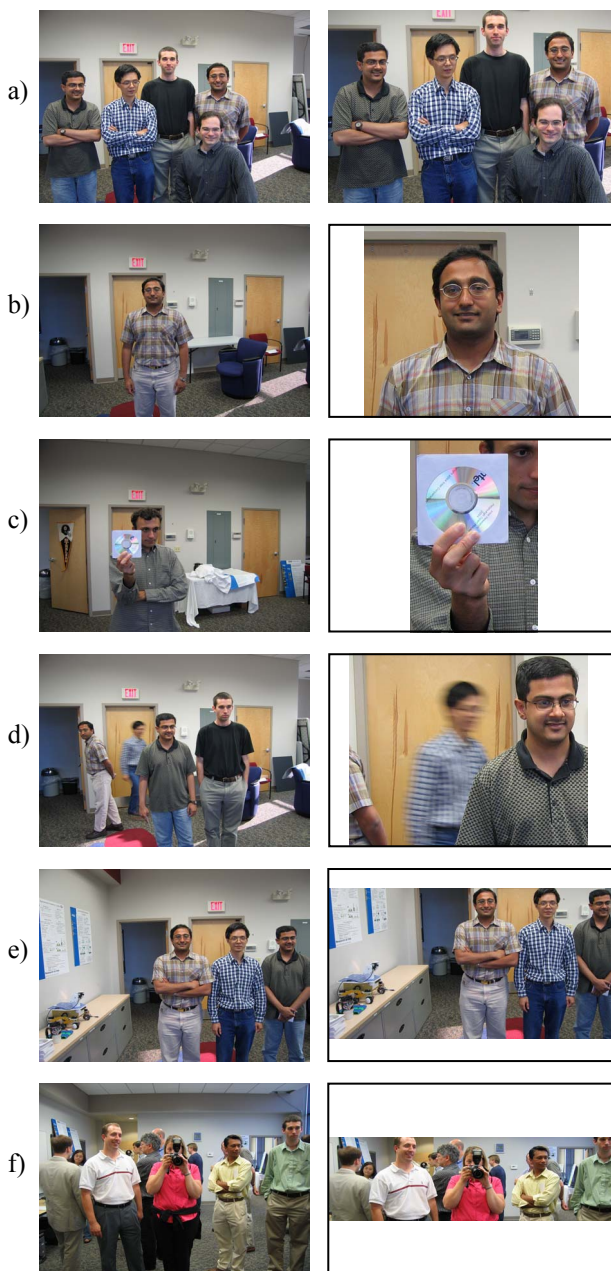


Fig. 7: Examples of successful (a,b,c) and unsuccessful (d,e,f) photographs taken by the Robot Photographer
Images are shown before and after automatic cropping

characteristics of posed group photos and applying optical flow and motion parallax techniques. The subject of a photo is identified, framed, and the image correspondingly cropped without resorting to the analysis of image content or an attempt to resolve the general image understanding problem. As a result, the Robot Photographer presented here works relatively well for both human and nonhuman subjects, and does not require heuristics or training specific to potential subjects.

While we do not expect the Robot Photographer to put professional photographers out of a job, it does achieve reasonable photographic performance. Further improvements in framing may be achieved by developing a hybrid approach that uses the motion-based techniques explored here in conjunction with the more traditional content-based techniques to more accurately identify the subject area across a wider range of scenes than with either type of algorithm alone. One potential picture quality refinement that is made possible by our technique is the selection of lens aperture based on subject depth range information to ensure sufficient depth of field, while simultaneously minimizing the shutter time to avoid motion blur of the subject. (Rather than simply relying on the arbitrary programmed exposure mode of the camera). Finally, the subject identification and framing techniques described here may also be more generally useful in scenarios with less restricted robot motion, such as in an exploration rover. Such a robot can certainly use our framing techniques to generate well-cropped images of features photographed, but may also be able to apply our techniques to the converse problem—positioning itself to ensure a subject is well-framed in the camera view.

ACKNOWLEDGMENTS

We wish to thank Emily Hamner, Eric Porter, Brian Dunlavey, Illah Nourbakhsh and the CMU Personal Rover Project for the loan of a Personal Exploration Rover for this work. We also thank Aroon Pahwa for providing robot control software for the PER.

REFERENCES

- [1] Z. Byers, M. Dixon, K. Goodier, C. Grimm, and W. Smart, "An Autonomous Robot Photographer," *Proceedings of the International Conference on Robots and Systems (IROS 2003)*, October 2003.
- [2] Z. Byers, M. Dixon, K. Goodier, W. Smart, and C. Grimm. "Say Cheese!: Experiences with a Robot Photographer," *Proceedings of the Fifteenth Innovative Applications of Artificial Intelligence Conference (IAAI-03)*, August 2003.
- [3] W. Smart, C. Grimm, M. Dixon, and Z. Byers. "(Not) Interacting with a Robot Photographer," *Proceedings of the AAAI Spring Symposium on Human Interaction with Autonomous Systems in Complex Environments*, March 2003.
- [4] Personal Rover Project: <http://www.cs.cmu.edu/~personalrover/>
- [5] J. Campbell, R. Sukthankar, I. Nourbakhsh. "Techniques for Evaluating Optical Flow in Extreme Terrain," *Proceedings of IROS 2004*.
- [6] Libptp2 package: <http://sourceforge.net/projects/libptp>
- [7] Ptpcanon package: <http://jalkapallo.org/superkolik/canon.html>
- [8] B. Lucas and T. Kanade. "An Iterative Image Registration Technique with an Application to Stereo Vision," *International Joint Conference on Artificial Intelligence*, pp 674-679, 1981.
- [9] J.-Y. Bouguet. "Pyramidal Implementation of the Lucas Kanade Feature Tracker," OpenCV Documentation, Intel Corporation, Microprocessor Research Labs, 1999.
- [10] OpenCV, <http://www.intel.com/research/mrl/research/opencv/>
- [11] J. Campbell, R. Sukthankar, I. Nourbakhsh, A. Pahwa. "A Robust Visual Odometry and Precipice Detection System Using Consumer-grade Monocular Vision," in press, *Proceedings of ICRA 2005*.