

The autonomous hidden camera crew

Timothy Callemein*

timothy.callemein@kuleuven.be

Wiebe Van Ranst*

wiebe.vanranst@kuleuven.be

Toon Goedemé

toon.goedeme@kuleuven.be

EAVISE, KU LEUVEN, BELGIUM

Abstract

Reality TV shows that follow people in their day-to-day lives are not a new concept. However, the traditional methods used in the industry require a lot of manual labor and need the presence of at least one physical camera man. Because of this, the subjects tend to behave differently when they are aware of being recorded.

This paper will present an approach to follow people in their day-to-day lives, for long periods of time (months to years), while being as unobtrusive as possible. To do this, we use unmanned cinematographically-aware cameras hidden in people's houses. Our contribution in this paper is twofold: First, we create a system to limit the amount of recorded data by intelligently controlling a video switch matrix, in combination with a multi-channel recorder. Second, we create a virtual camera man by controlling a PTZ camera to automatically make cinematographically pleasing shots.

Throughout this paper, we worked closely with a real camera crew. This enabled us to compare the results of our system to the work of trained professionals.

1 Introduction

Disturbing events occur in each household, both negative such as death, divorce, and disease, as well as positive ones like birth, adoption, and marriage. An independent video production house wants to make a reality TV show that will get families on TV in these situations, hoping to yield emotional and socially inspiring footage. Take for example first-time parents questioning how other couples manage to get up at night to care for their newborn child.

Recording a household with cameras is not groundbreaking in itself, various TV programs have already explored this concept. However, the production house desired an innovative system that was able to automatically record using on side hardware without the presence of a human camera crew. In this paper, we describe the solution we developed for this, composed of a system to decide which camera will be recording (section 4), and a controller for the PTZ to take cinematographic medium shots (as described in section 5). Remark that the video streams produced by our system form the input for a human film editor. In section 6 we present the results, both quantitatively and qualitatively by validating it using an experienced camera crew.

*Timothy Callemein and Wiebe Van Ranst contributed equally to this paper.




Figure 1. Hardware scheme

2 Related work

There are different camera network techniques available like [9, 10], that mainly focus on security applications, mostly involving far away cameras capturing an overview of an environment. In this paper however, cameras will follow cinematographic rules that will be positioned at eye height and fairly close to what they are capturing. The presented system will decide on the spot whether or not to record, making it critical to be fast and reliable, as opposed to surveillance video that is permanently recorded and stored for future analysis. Other techniques like [2, 4] that follow cinematographic rules edit the video streams to produce one video montage or only focus on one room. Our system will select multiple cameras placed in multiple rooms and record them minimizing data collection providing only useful material that still needs manual editing. This relatively new research topic is being actively researched [1].

Automatically controlling a pan-tilt-zoom (PTZ) camera to follow people is not a new concept. In [7] a system is proposed for lecture recording that also takes into account some cinematographic rules. The lecturer is tracked in the PTZ video (no overview camera is used) using a combination of a detector and a tracker, they use a PID loop to continuously keep the lecturer on a fixed position in the frame. While such a system works well for lecture recording, recording people in their house is much more challenging. Our system is able to create shots of multiple people, conform to different cinematographic rules and take pleasing static shots.

3 Hardware setup

Figure 1 illustrates the hardware setup comprised of multiple Full HD cameras to be recorded. Although the recorder is capable of capturing eight channels simultaneously, we expanded this number by adding a matrix allowing up to 20 cameras in total.

The chosen ProRes LT 422 codec captures at an av-




Figure 2. Top view with cameras, buckets and zones adding to the buckets.

erage of 102 Mbps, amounting up to 1 TB per continuous camera stream per day. The proposed camera selection method will reduce the needed storage, extending the recording time greatly, and reducing the search space of the human editor.

Two types of cameras are used, static and PTZ cameras, both with a dedicated SDI recording interface and IP interface for image processing and PTZ control. We add an overview camera to each PTZ, to detect actors and steer the PTZ as described in section 5.

We develop a semi-autonomous calibration procedure to calibrate the PTZ control parameters to the field-of-view of the overview camera. After a rough manual calibration, for each position SURF feature matching is used to correct the local parallax error and to fine align the FoV of the PTZ and the desired canvas cutout from the overview camera.

Our server processes the 20 camera streams in low-resolution through the network, allowing image based decisions that will select the cameras to record (described in section 4) and control the PTZ (section 5).

4 Camera selection

One of the camera crew’s tasks is deciding when to record, reducing large amounts of useless video material. This requires an algorithm that is able to anticipate the whereabouts of the residents, and their movements between rooms.

To make the record decision automatically, we need to understand what a camera crew uses as criteria for recording cinematographic scenes. In a reality TV show that follows the residents, the presence of a person is the primary trigger. In most cases, a script is used to determine when an actor is on set. In reality TV this is not possible, we have to use the entering or exiting of a resident as cue to start recording, which are shots that are also usable in editing. Movie professionals therefore prefer to have a pre-roll: shortly before a subject enters a room, the camera should be switched on.

Figure 2 shows a possible setup, three rooms and four partially overlapping cameras. On figure 2a two zones, blue and red, are doors connecting the three rooms. The red zone is visible on C_1 , C_2 , and C_3 . Every room also contains a virtual “bucket” that holds the room’s interest level (IL). The door in red, will raise IL of the buckets A and B. When a person is present in the red zone of a camera, the connected buckets will fill with water. Then, when a threshold is reached, we start recording the cameras connected to the bucket. This enables cameras that are in a different room, to trigger a room before entering.

The cameras project the 3D environment on a 2D image, this means a person can be inside a zone, yet far away from a door opening, leading to false positives.

Yet, as shown in figure 2b, multiple cameras will raise the IL faster, confirming the 3D location from multiple perspectives, providing more speed and reliability when triggering the recording. This combined with different weights per zone, limits false positives and provides fine-tuning capabilities. The buckets also contain a “hole”, that will lower the IL over time, greatly reducing the chances of accumulated low activity to trigger a false positive.

We use dedicated static cameras as well as the PTZ overview cameras to accurately monitor rooms inside a home. Since both of these cameras are in fixed positions, all movement in the video is caused by changes in the scene, meaning background subtraction (BGS) techniques work well.

In all of these cameras’ frames we define zones (eg. a door opening) that fill up accompanying buckets. We determine the fill-up rate of a bucket by multiplying the amount of foreground activity in a zone with preset weights that we assigned to each zone. This ensures that eg. a cat walking through a door has less influence than a person.

5 Autonomous PTZ steering

The role of the PTZ camera in our setup is to replace a traditional camera man without the need for any human intervention. By adapting the three degrees of freedom of a PTZ to all kinds of scenarios that take place in front of the camera, we are able to make shots that are a lot more interesting than those captured by a static camera. To implement such an “intelligent camera man” it is essential to have a good understanding of 1. what is currently happening in front of the camera and 2. given the knowledge of the current situation how to capture it in a cinematographically pleasing way. Note that, contrary to other automated PTZ systems like [7] our system’s main functionality is to make steady shots. For example, when someone is sitting on a couch watching TV, we do not continuously adjust the camera to keep the, sometimes slightly moving, person at a fixed position in the frame (often determined by the rule of thirds). Cinematographically speaking, such a “quivering shot” is not wanted. Instead, we wait for the person to become stationary, calculate what would be a good shot and then when we determine that the time is right, control the PTZ to make the shot.

5.1 Understanding the scene

In our application, the objective of our virtual camera man is to make shots of people. An obvious first step in creating such a system is to find out where people are located in the scene. To do this we compared different person detectors, like Deformable Parts Model (DPM) [5, 6] or Aggregate Channel Features (ACF) [3] combined with background subtraction and some scene constraints to speed up the algorithm. For accuracy (especially when dealing with occlusion) and because speed is not our main concern when creating static shots (real-time performance is not necessary) we choose to run a DPM upper-body detector. In section 6 we compare these detectors in more detail on our own dataset.

Apart from a person’s location and size, which we get from the person detector, we also need to know the orientation and the position of the head inside the detection. For this, we use the approach proposed by [8]. This technique uses the Viola and Jones [11] face detector with three separate models, one trained on the frontal views, and the other on the left and right profile views of a face dataset. The results are then combined to estimate the gaze of the detected person. Because this method does not rely on small-grained facial features, it is proven to also work on small faces with limited resolution.

Apart from this, we also managed to speed up the detector by combining it with background subtraction. When we detect movement in the foreground image we cut out bounding boxes around active areas. These areas are stored in a pool on which the detector is evaluated. The resulting bounding boxes of the detector (with some added margin) are then added to this pool to be evaluated again (together with the resulting bounding boxes of background subtraction) in the next image. Evaluating places where people were found in the previous image, prevents people who are stationary, ie. invisible to the background subtraction algorithm, from disappearing.

5.2 Making cinematographically pleasing shots

The detector provides us with the necessary information to produce cinematographic shots: the location, size, position and general gaze direction for each person, detected in the overview camera. We give priority to keeping the maximum amount of people in one shot, over making medium shots with the rule of thirds. This means that the distance between the outer detections will determine the width and the height by the 16:9 ratio of the shots. To provide enough space between the outer detections and the border, we add 15% to the width. Vertically, we place the estimated eye position of the highest detection to comply with the rule of thirds.

When there is only one detection found the height and aspect ratio of the detection is used as reference to determine the size of the frame. The position of the eyes is again used to determine the position of the frame on the y-axis. While the general gaze will now determine the position on the x-axis. When a person is looking to the right, we leave space to the right, when he is looking to the left, we provide space to the left. When there is no general gaze available, we choose to place the detection in the center. This is not optimal from a cinematographic perspective, but results in more usable shots than choosing one side as default. Some examples of the resulting shots are illustrated in figure 6.

5.3 Switching between different shots

As mentioned before, in our application, a continuously moving “quivering” camera is not desirable, we want to make steady shots. To make sure our camera makes steady shots and to prevent it from constantly alternating between different shots, we adjust the camera only if following conditions are met:

- The proposed shot is considerably different (determined by position and size) from the current.




Figure 3. Camera selection results of three buckets during a 10 minute scenario.




Figure 4. Comparison of the time complexity of different detection algorithms on our test footage.

- The new shot is steady, meaning people in it are only moving slightly or not at all.
- The old shot was kept a bit longer (2s in our case), to allow subjects to leave the scene, and has a duration of at least 6s.

If all these conditions are met, we map the proposed canvas from overview camera coordinates to PTZ control parameters.

6 Experimental results

In this section, we will report on quantitative results of the three main components of the proposed system (the camera selection, person detector and the cinematographical canvas selection method). However, as the proposed system is currently being used in a grand scale reality TV video production (8 families filmed 24/7 over 12 months), the very evaluation of the system will be in the success of that TV show.

6.1 Camera Selection

To test the camera selection algorithm we played a scenario inside the test environment, which included movements between different rooms. Figure 3 shows the analysis of three buckets being triggered, triggering recordings, during the 10 minute scenario, which is sampled every 10s and analyzed by comparing the actual output with the expected recording, as given by a real camera man. Since it’s important to capture everything (and not lose any action) we set the sensitivity quite high, explaining the presence of multiple false positives (FPs) (triggering while no activity) and also the lack of false negatives (FNs) (activity but no or to late triggering of recordings). These results show that the buckets trigger the connected cameras on time, while not missing any vital video information. The accuracy of the tests by room are: hallway 100%, living room 87% and kitchen 95%. In this experiment, we saved 43.33% of disk space opposed to recording the full 10 minutes and recorded only 8% data overhead that contains no usable video data.

6.2 Person detector

In figure 4 we compare the time complexity of the different upper body detection algorithms on a test dataset that was recorded in a real setting (inside

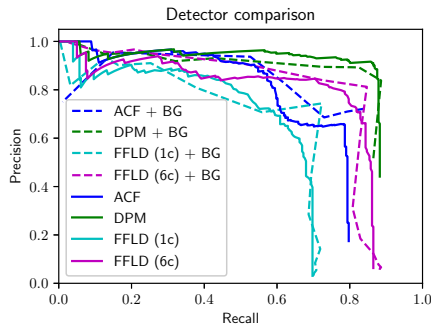


Figure 5. Detection performance on overview footage from a typical home.

a typical home). We compare the algorithm with background subtraction enabled (“+ BG” in figure 4). While using background subtraction has a much faster average execution time, there are cases where the entire image needs to be searched causing a big performance hit, for instance when a person walks in front of the camera.

Figure 5 compares the performance of these detectors on the same dataset. We annotated a total of 105 frames with a total video duration of about two minutes. Some samples of this dataset can also be seen in figure 6. It is clear that the DPM upper body detectors (ffld_6c and dpm) [6, 5] have the best performance (at the cost of being more computationally intensive), compared to the rigid ACF [3] detector which is not as accurate but much faster. Limiting the search space by using background subtraction seems to have little impact on detector performance. Adding more people to the scene will only impact detection speed.

6.3 PTZ canvas

Figure 6 shows the overview image with the detections and the resulting red frame that will be used to control the PTZ. As can be seen, the resulting canvases are designed to be cinematographically pleasing. To ensure this, we repeated the same 10 minute scenario 10 times, and continuously let the resulting canvases evaluate by a group of trained camera men. Some video material of these tests can be seen at <http://goo.gl/WIS0P4>, containing video footage that is selected to be recorded (the multicam video) and multiple videos with detections and proposed canvases.

7 Conclusion

In this paper, we proposed a system to automatically film indoor scenarios with multiple static and PTZ cameras, while complying to cinematographic rules. The first task of an autonomous camera crew is choosing whether to record or not. The test results show that our proposed bucket system is able to record all useful video streams with little overhead, decreasing the needed recording storage and search space for the human editor. The next task was using a PTZ that is able to take medium shots of people detected in an overview camera. Results show that the PTZ is taking multiple cinematographically pleasing static shots of single and multiple persons based on the overview




Figure 6. Proposed canvases (red) based on people detections (blue) controlling the PTZ. Video illustration on <http://goo.gl/WIS0P4>.

camera info. Remark that the system we have developed is currently being deployed in a real reality TV video production, providing us with more data to be validated.

Acknowledgments

This work is supported by video production house Geronimo and the KU Leuven Cametron project.

References

- [1] J. CHEN AND P. CARR, *Autonomous camera systems: A survey*, in AAAI-14 Workshop on Intelligent Cinematography and Editing, 2014.
- [2] F. DANİYAL AND A. CAVALLARO, *Multi-camera scheduling for video production*, in Visual Media Production (CVMP), 2011 Conference for, IEEE, 2011, pp. 11–20.
- [3] P. DOLLÁR, R. APPEL, S. BELONGIE, AND P. PERONA, *Fast feature pyramids for object detection*, TPAMI, 36 (2014), pp. 1532–1545.
- [4] P. DOUBEK, I. GEYS, T. SVOBODA, AND L. VAN GOOL, *Cinematographic rules applied to a camera network*, in Omnivis 2004, Prague, 2004, pp. 17–29.
- [5] C. DUBOUT AND F. FLEURET, *Exact acceleration of linear object detectors*, in European Conference on Computer Vision, Springer, 2012, pp. 301–311.
- [6] P. F. FELZENSZWALB, R. B. GIRSHICK, AND D. MCALLESTER, *Cascade object detection with deformable part models*, in CVPR, IEEE, 2010, pp. 2241–2248.
- [7] D. HULENS, T. GOEDEMÉ, AND T. RUMES, *Autonomous lecture recording with a ptz camera while complying with cinematographic rules*, in CRV, 2014.
- [8] D. HULENS, K. VAN BEECK, AND T. GOEDEMÉ, *Fast and accurate face orientation measurement in low resolution images on embedded hardware*, in VISIGRAPP, vol. 4, 2016, pp. 538–544.
- [9] L. LIAO, D. FOX, AND H. KAUTZ, *Blocation-based activity recognition using relational markov networks*, in *proc. 19th int.*, in JCAI, Edinburgh, Scotland, 2005, pp. 773–778.
- [10] B. SONG, A. T. KAMAL, C. SOTO, C. DING, J. A. FARRELL, AND A. K. ROY-CHOWDHURY, *Tracking and activity recognition through consensus in distributed camera networks*, TIP, 19 (2010), pp. 2564–2579.
- [11] P. VIOLA AND M. JONES, *Rapid object detection using a boosted cascade of simple features*, in CVPR, vol. 1, IEEE, 2001, pp. I–511.