

Speaker-Following Video Subtitles

YONGTAO HU, The University of Hong Kong

JAN KAUTZ, University College London

YIZHOU YU and WENPING WANG, The University of Hong Kong

We propose a new method for improving the presentation of subtitles in video (e.g., TV and movies). With conventional subtitles, the viewer has to constantly look away from the main viewing area to read the subtitles at the bottom of the screen, which disrupts the viewing experience and causes unnecessary eyestrain. Our method places on-screen subtitles next to the respective speakers to allow the viewer to follow the visual content while simultaneously reading the subtitles. We use novel identification algorithms to detect the speakers based on audio and visual information. Then the placement of the subtitles is determined using global optimization. A comprehensive usability study indicated that our subtitle placement method outperformed both conventional fixed-position subtitling and another previous dynamic subtitling method in terms of enhancing the overall viewing experience and reducing eyestrain.

Categories and Subject Descriptors: H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—*Evaluation/methodology*; C.4 [Performance of Systems]: Design Studies; H.1.2 [Models and Principles]: User/Machine Systems—*Human factors*

General Terms: Algorithms, Design, Performance

Additional Key Words and Phrases: Speaker detection, speaker-following subtitle placement, video viewing experience

ACM Reference Format:

Yongtao Hu, Jan Kautz, Yizhou Yu, and Wenping Wang. 2014. Speaker-following video subtitles. *ACM Trans. Multimedia Comput. Commun. Appl.* 11, 2, Article 32 (December 2014), 17 pages.

DOI: <http://dx.doi.org/10.1145/2632111>

1. INTRODUCTION

Subtitles are necessary for television programs, movies and other visual media for people with hearing impairments to help them understand and follow the dialog. Translated subtitles are necessary for media in a foreign language that is not dubbed to help viewers understand the spoken dialog. Subtitling is also useful for learning foreign languages.

Conventionally, the position of the subtitles is in a fixed location such as at the bottom of the screen. Humans can see objects within their field of vision but can only read text clearly in a narrow vision span. This vision span is an angular span (vertically and horizontally) of approximately 6 degrees of arc, which yields a region with a diameter of 5.23cm when viewed from 50cm away (Figure 1) [Rayner 1975; McConkie et al. 1989; Just and Carpenter 1987; Wikipedia 2012]. The conventional location of subtitles at the bottom of the screen means that to follow both the facial expression of speakers and the subtitles, the viewer has to constantly move their eye gaze between the main viewing area and the bottom of the screen, leading to a high level of eyestrain.

Author's address: Y. Hu; email: herohuyongtao@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2014 ACM 1551-6857/2014/12-ART32 \$15.00

DOI: <http://dx.doi.org/10.1145/2632111>




Fig. 1. Limited vision span of human eyes. The human vision span is an angular span of 6 degrees of arc both vertically and horizontally, which yields a region with a diameter of 5.23cm when viewed from 50cm away.



Fig. 2. Previous work to improve accessibility and understanding of conversations through word balloon layout in comics (a, b), and subtitle placement in videos (c). All images with permission of the authors of corresponding papers.

Considerable effort has been made to enable the viewer to understand conversations through better presentation of the spoken dialog. In comics, the word balloon construction and layout of the graphics are techniques used to achieve this goal [Kurlander et al. 1996; Chun et al. 2006]. Word balloon layouts in comics are used to help readers better visualize the written dialog, and the layout is carefully arranged to ensure that the flow of the story is clear (Figure 2(a) and 2(b)). The previous work by Hong et al. [2010] is the first and only research that extends these concepts to enhance the accessibility of videos by placing on screen subtitles next to the respective speaker (Figure 2(c)).

In this article, we will analyze the placement of subtitles so the viewer can comfortably follow the video content and subtitles simultaneously, which enhances the viewing experience. We will improve on the work by Hong et al. [2010] in two main aspects. We present a new speaker detection algorithm to accurately detect the speakers based on visual and audio information. An efficient optimization algorithm is then used to place the subtitle based on a number of factors. The framework of our approach is shown in Figure 3.

To summarize, the contributions of this article are:

- the development of an automatic end-to-end system for subtitle placement for TV/movies;




Fig. 3. Our framework consists of two main components: 1) speaker detection; and 2) smart subtitle placement.

- a new algorithm for effectively detecting speakers based on visual and audio information;
- an efficient optimization algorithm for determining appropriate locations to place subtitles.

The rest of the article is organized as follows. Related works are discussed in Section 2. The problem formulation and the whole framework are introduced in Section 3. Experiments results are presented in Section 4, followed by conclusions and future works in Section 5.

2. RELATED WORK

The work of Hong et al. [2010] is the first and only previous research to study how dynamically placed subtitles can enhance video accessibility for the hearing impaired. The speaker is determined by examining lip motion features and then subtitles are placed in a nonsalient region (based on a saliency map) around the speaker. Their usability study showed that their method effectively improved video accessibility for viewers with hearing impairments. We have extended this work to improve video subtitling and enhance viewing experience. First, we developed a new algorithm for improving speaker detection accuracy. The method is based on both visual and audio information rather than using only lip motion because other people in the scene may also be moving their lips. Second, the subtitles are positioned using an optimization procedure that takes several factors into consideration, including absence of speakers, cross-frame coherence, and screen layout. The combined use of these factors has proven to be more robust than using image saliency alone.

Speaker detection in video is an active research topic. Speaker diarization is the problem of determining “who spoke when” based on audio signals [Anguera Miro et al. 2012]. The problem of speech recognition based on visual information has been studied by Gordan et al. [2002] and Saenko et al. [2004, 2005]. Visual and audio data are often used together in speaker localization [Nock et al. 2003; Potamianos et al. 2003]. Generally speaking, methods based on training data are more accurate but are more time-consuming. Out of efficiency consideration, we developed a method without involving a training process but still achieved robust detection results because of our use of several novel features and cues.

Subtitle placement has been previously investigated by Kurlander et al. [1996] and Chun et al. [2006] in their work on word balloon layouts in comics. Their methods are focused on handling single frames and so cannot easily be extended to video subtitles because they cannot ensure cross-frame coherence. In contrast, our approach employs a comprehensive optimization procedure to ensure cross-frame coherence to improve video subtitle presentation.

81	82
00:03:00,400 --> 00:03:02,400	00:03:02,402 --> 00:03:04,819
It's part of my sword collection.	- Do you have a sword collection? - No.

Fig. 4. Two examples of subtitle segments. A subtitle segment has three parts: the segment index (the first line), the timing information of the segment (the second line), and the words spoken in this segment (the third line). The first subtitle segment is for one speaker and the second subtitle segment for two speakers.

3. METHODOLOGY

3.1. Problem Formulation

Given a video (TV, movie, etc.) and its subtitle (caption) file as input, our goal is to generate the same video but with the subtitles positioned next to their corresponding speakers to provide better viewing experience. Our method consists of two main components: 1) speaker detection; and 2) optimization of subtitle positions. The workflow of the method is shown in Figure 3.

3.2. Preliminaries

The subtitle file is a text file in SRT format consisting of subtitle segments that contain the spoken lines and timing information. Time information provides only an approximate time interval during which the subtitle are shown on the video screen. Since this interval is always longer than the speaking time with extra time added at the beginning and the end. The input video consists of video segments corresponding the timing information in the subtitle file. Those video segments that have corresponding subtitle segments are called *speaking video segments*, while the others called *nonspeaking video segments*. Figure 4 shows an example of two consecutive subtitle segments. A subtitle segment may contain lines from more than one speaker, as in the second example. The first example in Figure 4 corresponds to one speaking video segment, while the second example to two speaking video segments for two speakers.

We will explain next how to process all the speaking video segments for speaker detection and optimization of subtitle placement.

3.3. Speaker Detection

For each speaking video segment, we perform the speaker detection procedure to detect the speakers. The workflow for speaker detection is shown in light blue in Figure 3. Our strategy is to initially detect the faces and combine them to form face tracklets. We examine all the face tracklets to determine the speaker for a given subtitle segment. The details of each step are discussed in the following text.

3.3.1. Face Tracking. We first detect faces using the OpenCV frontal + profile Viola-Jones face detector [Viola and Jones 2004]. Then these detected faces are linked using a low-level association approach to obtain face tracklets, following Kuo et al. [2010]. Regions with similar positions, sizes and appearances are linked. These low level tracklets are further linked into final face tracklets if their size, appearance and coherent direction of movement are very similar to each other.

Face detections are not directly used for face tracking because matching the appearances of faces is extremely challenging due to similar colors, expressions, poses of different faces, in addition to lighting changes and motion blur [Zhang et al. 2003]. Therefore, the clothing appearance has been used as additional cues for matching the same face [Jaffré et al. 2004]. Previous works show that face tracking in certain conditions, where matching is not possible based on face appearance alone, can be



Fig. 5. Clothing appearance used in face tracking. Red rectangles show the detected faces and yellow rectangles show the corresponding clothing appearance. Screenshot with permission of the producer of movie “The Man from Earth”.

overcome by matching clothing appearances [Everingham et al. 2006]. This improvement is mainly due to the richer texture variety of clothing compared to faces. Hence, we incorporated clothing appearance in our face tracking algorithm and found that this improved the performance significantly. The bounding boxes of clothing are extracted according to the localization and the scale of the detected faces. We estimate the clothing appearance in an area under the face at a distance of $0.2 \times$ face height. The size of this rectangular clothing area is proportional to the face area with a width $2d$ and height $1.5d$, where d is the width of the face. Examples are given in Figure 5. We chose these coefficients from the best fitting of the box in learning images.

3.3.2. Speaker Detection. Lip motion is the most widely used feature for speaker detection [Everingham et al. 2006; Hong et al. 2010]. However, the lip motion feature alone is not robust enough for speaker detection because the detection of a mouth region is often not accurate. We have observed that in TV and movies the positions and duration of the speakers are highly distinguishable from non-speakers. In addition to lip motion, we introduce a center contribution feature and length consistency feature for better speaker detection. We further consider the audio-visual synchrony feature to take into account the close relationship between audio and visual information to improve the robustness of our method. Based on these observation, in the presence of multiple speaker candidates, we detect the true speaker by examining the following four features: 1) *lip motion*, 2) *center contribution*, 3) *length consistency* and 4) *audio-visual synchrony*. The detailed procedures are explained next.

3.3.2.1. Lip Motion. A speaker often has more significant lip motion than nonspeakers. Similar to Everingham et al. [2006], we can detect the speaker from speaker candidates by detecting those faces with significant lip motion. For each face tracklet, we first use a facial landmark detector [Uricár et al. 2012] to detect mouth corners. The mouth region can then easily be predicted based on the mouth corners. The mean squared difference (MSD) of the pixel values in mouth region is then computed between the current and previous frames. We compute the average of MSDs as the lip motion feature as

$$\mathcal{MSD} = \frac{1}{N-1} \left[\sum_{i=1}^{N-1} \mathcal{MSD}_{(M_i, M_{i+1})} \right], \quad (1)$$

where N is the length (i.e. # of frames) of the speaker tracklets, M_i the mouth region of the speaker in the i -th frame.

3.3.2.2. Center Contribution. It is observed that, in most TV and movies, a speaker is more likely than nonspeakers to be located towards the center of the screen. To

leverage this observation, we introduce the *center contribution* (CC) feature as

$$CC = \frac{1}{N} \left[\sum_{i=1}^N CC_i \right], \quad \text{with } CC_i = 100 \times \left[1 - \frac{d(P_i, P_c)}{d(\mathcal{O}, P_c)} \right], \quad (2)$$

where P_i is the center of the speaker's face in i -th frame, and \mathcal{O} and P_c are the origin and center of the image plane, respectively. Here $d(\cdot, \cdot)$ is the distance function between two points in Euclidean space.

3.3.2.3. Length Consistency. The accuracy of speaker detection can be further improved by considering the consistence between the length of a candidate face tracklet and the length of speaking time of a speaking subtitle segment. We call this *length consistency*. Similar to MSD and CC, the candidate face tracklet with a higher LC is more likely to be the true speaker. We compute the length consistency (LC) as follows

$$LC = \frac{1}{|L - L_{std}|}, \quad \text{with} \quad \begin{cases} L_{std} = \frac{L_{words} \cdot F_{video}}{V_{speaking}}, \\ V_{speaking} = \frac{L_{total_words}}{T}, \end{cases} \quad (3)$$

where L is the length of the candidate face tracklet, L_{words} the number of the spoken words in the current subtitle segment, and L_{total_words} the total number of words in the input video, T the total speaking time in the input video, F_{video} the frame rate of input video, and $V_{speaking}$ the average speaking speed in the input video (i.e., the number of words spoken in unit time).

3.3.2.4. Audio-Visual Synchrony. Previous work [Driver 1996; Wallace et al. 2004] has demonstrated that audio cues and video cues can be combined to enhance the understanding of an environment. For example, sounds appear to correspond to motion synchronous with acoustic stimuli. Audio-visual (AV) synchrony has been used to resolve speaker localization [Monaci 2011] by computing the synchronization score of audio with the lower half of the faces, with audio-visual cooccurrence measured by the *synchronization score* \mathcal{AV} of time slot $\Delta = [T_1, T_2]$ which is computed as

$$\mathcal{AV} = \langle y_a(t), y_v(t) \rangle, \quad \text{with } t \in \Delta, \quad (4)$$

where y_a , y_v are the audio and visual feature vectors and $\langle \cdot, \cdot \rangle$ indicates the scalar product between the vectors.

Speaker detection that is solely based on lip motion [Hong et al. 2010] relies too much on the accuracy of the mouth region detection and therefore may fail for complex scenes. Our method includes *center contribution* and *length consistency* features, and we further improve speaker detection by examining *audio-visual synchrony*. Moreover, unlike Monaci [2011], we use extended facial landmarks [Uricár et al. 2012] for better prediction of the motion region instead of using only the lower half of the face when examining audio-visual synchrony. This will help to filter out any motion disturbance in the background, especially when speakers or cameras are moving. Our method also compensates for the errors due to poor face detection and facial landmarks (see Figure 6).

For better accuracy and efficiency, we chain the four features mentioned before in a cascade in the following order: lip motion, center contribution, length consistency and audio-visual synchrony (see Figure 7). The process can be seen as a degenerate decision tree [Viola and Jones 2004]. The design of this cascade structure is based on the observation that the majority of the candidate speakers are nonspeakers and only candidate speakers that satisfy the current constraint pass to the next level. We put lip motion first in the cascade because it can effectively reject most nonspeakers. Center



Fig. 6. Improving audio-visual synchrony through better motion region prediction. Left two: motion region (light pink area) used by Monaci [2011]. Right two: better motion region (light yellow area) prediction based on facial landmarks (red points) from our method. Screenshots with permission of the producer of movie *The Man from Earth*.




Fig. 7. Schematic depiction of our cascade speaker detection method. Only candidate speakers that pass the current test will go through to the next level.

contribution and length consistency are put second and third in the cascade as they can be evaluated very quickly. Audio-visual synchrony is placed last in the chain because this evaluation is more time-consuming. Experiments have also shown that a cascade in this order gives the best results. Detailed implementation is shown in Algorithm 1.¹

Figure 11 shows the performance improvement in speaker detection accuracy.² The left figure shows the effect of adding the center contribution, length consistency and audio-visual synchrony (with better motion prediction) besides *lip motion*, and the right figure shows our speaker detection accuracy compared with previous methods [Monaci 2011; Hong et al. 2010].³ Overall, our speaker detection method outperformed these two methods.

3.4. Subtitle Placement

Generally speaking, our goal is to place subtitles close to their corresponding speakers. There are several considerations in attaining this goal: (1) In each frame, the subtitle should be close enough to its speaker so not to cause any confusion that it is spoken by another person; (2) Across consecutive frames, the overall distance between all subtitle placements should be small to reduce eyestrain; (3) Subtitles should not be placed near screen boundaries so as not to detract the viewer from the central viewing area; (4) Subtitles should not occlude any important visual contents (e.g., faces).

¹We apply $\theta_1 = 20$, $\theta_2 = 2.5$, $\theta_3 = 2$, $\theta_4 = 0.1$, $\theta_5 = 2$ throughout.

²Speaker detection accuracy is computed in terms of video segments. It is considered to be correct for those video segments where the subtitle is put in the default position (i.e., the bottom of the screen) when no speaker is detected.

³Hong et al. [2010] used only the *lip motion* feature (i.e., MSD) for speaker detection.

ALGORITHM 1: Speaker detection algorithm.**Input** : N face tracklets, each of which stands for a speaker candidate.**Output:** The speaker (tracklet) or no speaker.Compute MSD feature for each face tracklet.Delete all tracklets whose $MSD < \theta_1$ (assume remaining N_1 tracklets after this).**if** $N_1 \geq 2$ **then** Delete all tracklets whose $MSD * \theta_2 < \max_{MSD}$ (assume remaining N_2 tracklets after this). **if** $N_2 \geq 2$ **then** Compute CC feature for each face tracklet. Delete all tracklets whose $CC * \theta_3 < \max_{CC}$ (assume remaining N_3 tracklets after this). **if** $N_3 \geq 2$ **then** Compute LC feature for each face tracklet. **if** $\max_{LC} - \text{second max}_{LC} > \theta_4$ **then return** the tracklet with \max_{LC} .

;

else

Compute motion region based on facial landmarks for each tracklet.

 Compute AV feature for each face tracklet **if** $\max_{AV} > \theta_5$ **then return** the tracklet with \max_{AV} .

;

else return no speaker. ; **else if** $N_3 == 1$ **then return** current tracklet. ; **else return** no speaker. ; **else if** $N_2 == 1$ **then return** current tracklet. ; **else return** no speaker. ;**else if** $N_1 == 1$ **then return** current tracklet. ;**else return** no speaker. ;

In this section, we will present an optimization framework (detailed in Algorithm 2) for subtitle placement. First, we will explain some preprocessing steps before performing subtitle placement. The following operations are needed to process the input video file and the subtitle file.

- (1) Split a speaking video segment when the speaker moves significantly during a subtitle segment;
- (2) Split a subtitle segment if there is a major shot change in the corresponding video segment;
- (3) Split a subtitle segment when it corresponds to multiple speakers; and
- (4) Refine the display time of a subtitle from the provided timing information.

ALGORITHM 2: Subtitle placement algorithm.**Input** : Speaker detection result (i.e., face tracklets for each subtitle segment).**Output:** Positions for subtitle placement at each subtitle segment.**foreach** *subtitle segment* **do** Determine the proper rectangle to hold the subtitle.⁴

Compute candidate positions for placing the subtitle.

Compute the optimal positions of subtitle based on candidate positions.

return the position of each subtitle segment.

⁴Similar to Hong et al. [2010], we use a rectangle to bound the subtitle and an arrow to the speaker to avoid confusion. The dimensions of the bounding rectangle of the subtitle can easily be determined by processing the subtitle text based on its length and font.



Fig. 8. Candidate positions for the placement of the subtitle around the speaker's face. Speaker (yellow rectangle), nonspeaker (red rectangle), candidate subtitle positions (yellow circles, 8 positions in our experiment). Screenshot with permission of the producer of movie *The Man from Earth*.

3.4.1. Computing Candidate Subtitle Positions. Any position close to the speaker may be a good subtitle position. In order to reduce the search space for better efficiency, we only consider the following eight candidate positions around the speaker's face: above left, above, above right, below left, below, below right, left and right, which are indicated by the yellow dots in Figure 8). Note that the exact locations of these eight candidate positions are computed based on the speaker's location and size of the speaker's face and the length and font size of the subtitle.

3.4.2. Subtitle Position Optimization. Subtitle position optimization is the key step in Algorithm 2. We take into account layout information to obtain the best position to place the subtitle. Here, we assume that each subtitle will be assigned a fixed position.

3.4.2.1. Local optimization. Locally, we require the subtitle to be close to the speaker and far away from nonspeakers. Let P denote the optimal subtitle position of the current subtitle segment, P_S the speaker's position, and P_{NS_k} ($k = 1, 2, \dots, K$) the positions of the K nonspeakers in the scene. The energy function of P due to local optimization is defined as

$$E_{local} = d_{(P, P_S)} - \sum_{k=1}^K d_{(P, P_{NS_k})}. \quad (5)$$

3.4.2.2. Position consistency over time. We require that the positions of consecutive subtitles do not change their positions too much across consecutive frames in order to alleviate eye-strain. To this end, we constrain subtitles of subsequent speaking video segments to follow the positions of the preceding subtitles. Let P_{pre_opt} denote the optimized position of the preceding subtitle. Then the energy function due to this consideration is defined as

$$E_{global} = d_{(P, P_{pre_opt})}. \quad (6)$$

3.4.2.3. Layout with respect to screen boundary. Subtitles should not be placed at the screen boundaries in order to allow the viewer to focus more on the central viewing part of the screen. This preference is reflected in following energy function:

$$E_{layout} = d_{(P, boundary)}. \quad (7)$$

Now, combining these three energy terms, we defined the following total energy function.

$$\min E = w_1 \cdot E_{local} + w_2 \cdot E_{global} + w_3 \cdot E_{layout}. \quad (8)$$



Fig. 9. Results of shot change detections. Each row shows a video segment. The first frame of the shot change is marked by a red rectangle with a yellow shade. There is one shot change in the first video segment and two shot changes in the second segment. Screenshots with permission of the producer of movie “The Man from Earth”.

We use $w_1 = 1.0$, $w_2 = 1.0$, $w_3 = -1.0$ for all our experiments. We minimize this function to compute the optimized position of the current subtitle.

3.5. Some Further Considerations

3.5.1. Advanced Video Segment Splitting. When the speaker moves around during a subtitle segment, we would still like to have the subtitle tagged to the speaker. However, rather than letting the subtitle float with the speaker, we split both the subtitle segment and its corresponding speaking video segment into shorter segments. Such subtitle segments are first identified, based on the speaker tracklet, and are split into shorter subtitle segments. Then the positions of these shorter subtitle segments are determined using the same procedure (Algorithm 2) as for the normal subtitles, so that the viewer can better follow the moving speaker together with the subtitles. Note that, the position of each short subtitle is fixed for each small video segment. One example is shown in Figure 10(a).

3.5.2. Shot Change Handling. A significant shot change in TV/movies may make a speaker appear or disappear. In this case, subtitle positions need to change accordingly. For example, it should be placed at a default position (such as the bottom of the screen) if the speaker disappears after a shot change. We propose a simple shot change detector for video segments solely based on color histogram (see Algorithm 3⁵) instead of using shot change detectors based on complex models [Sethi and Patel 1995; Faernando et al. 2001; Dimou et al. 2005]. Our shot change detector proves to be accurate enough even for some very challenging TV series such as “Friends” and “The.Big.Bang.Theory”, as shown in Table I. Some results of detected shot changes are shown in Figure 9.

ALGORITHM 3: Shot change detector for video segments.

Input : A video segment.

Output: All shot change positions.

```

foreach pair of adjacent two frames do
  Compute the similarity of their color histogram as  $\delta$ .
  if  $\delta < threshold$  then
    Label the second frame as a shot change position.

```

return all these labeled positions.

Based on the shot change detection, video segments are further split into shorter video segments, each of which has no big significant shot changes and the subtitle

⁵We computed the correlation coefficient as the similarity of their RGB color histogram and applied a threshold of 0.99.

Table I. Performance of our shot change detection procedure. The ground truth of shot changes are manually labeled in our experiments

Video Input	True	Detected	False	Missed	Recall	Precision
Friends.S01E05	115	105	3	13	88.70%	97.14%
Friends.S01E13	103	96	2	9	91.26%	97.92%
Friends.S10E15	106	108	3	1	99.06%	97.22%
The.Big.Bang.Theory.S03E05	31	32	1	0	100%	96.88%

position in it can be kept relatively constant. After splitting, the video segment containing the scene that has the longest time overlap with the speaking video subtitle will be assigned the subtitle. The subtitle for all the other video segments will be set at a default position (e.g., the bottom of the screen). Some sample outputs are shown in Figure 10(b).

3.5.3. Splitting Multiple Speakers. There are often two or more speakers in a single subtitle segment, such as in the second example in Figure 4, where different speakers are denoted by a ‘-’ in these segments. In this case, we split such a subtitle segment into different parts corresponding to different speakers, using the separating character ‘-’. Accurate active speaking time for each speaker is determined based on all the face tracklets of the whole segment. Then, within each small time range, Algorithm 1 is used to find the specific speaker. One example is shown in Figure 10(c).

3.5.4. Refining the Speaking Time. The time information in the subtitle file is usually not the actual speaking time but only provides an approximate interval to show the subtitle on screen. The time interval is manually specified and is normally longer than the actual speaking time, with extra time padded at the beginning and the end. We refine the speaking time based on the span of the speaker tracklet, which is more reliable and coherent with the video content. Some examples are shown in Figures 10(a), (b), and (c).

4. EXPERIMENT RESULTS AND DISCUSSION

We tested our system on a variety of videos including those in which people speak very fast and move quickly. The videos were from the following TV series and movies: *Friends*, *The Big Bang Theory*, *Kramer vs. Kramer*, *Erin Brockovich*, *Up in the Air*, *The Man from Earth*, *Scent of a Woman*, and *Lions for Lambs*. All these video clips can be accessed at the website⁶. Some sample outputs (TV/movie screenshots) are shown in Figure 10.

4.1. Speaker Detection Accuracy

We have introduced four new features for speaker detection. Their contributions to improving detection accuracy are shown in Figure 11. The center contribution improved the speaker detection accuracy by 0.5%–2.5%, length consistency improved accuracy by 1.5%–4%, and audio-video synchrony with better motion prediction improved accuracy by 3.5%–6.5%. Taken together, our speaker detection method outperformed Monaci [2011] by 8.5%–20.5% and Hong et al. [2010] by 7.5%–17%.

4.2. Subtitle Placement

In the previous method [Hong et al. 2010], subtitles are positioned in the least salient region around the speaker based solely on a saliency map. This performs well only with a relatively simple background. For videos with a complex background such as

⁶<https://sites.google.com/site/smartsubtitles/>.

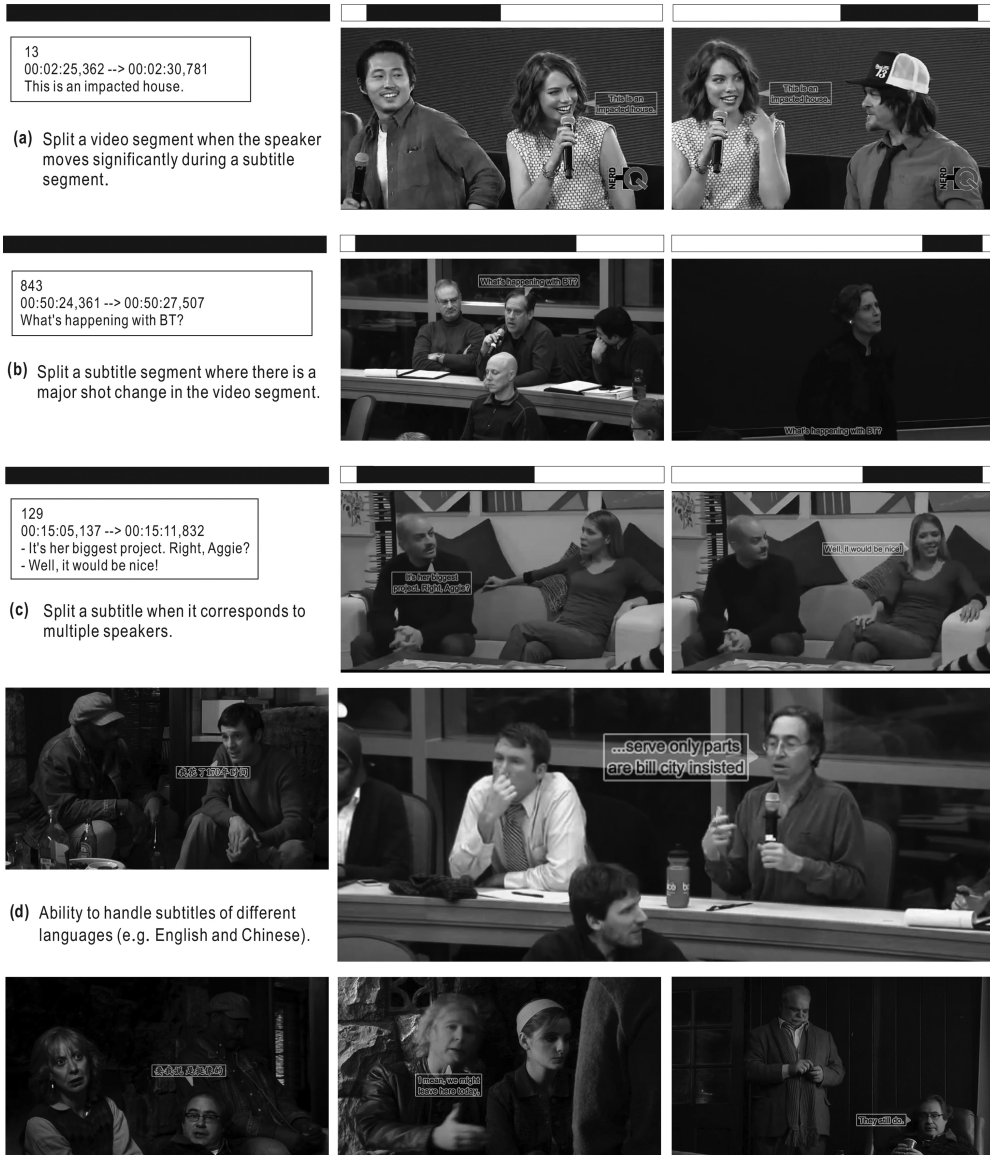


Fig. 10. Sample results of our system. In a) b) c), black/white bars are the active time (active in black, non-active in white). Screenshots of the first row with permission from YouTube video $v=v0y1rZXthmI$; the second row and the biggest one in the fourth row with permission from YouTube video $v=7R4xhwy-1oI$; the third row of with permission from YouTube video $v=CXP6KU4urgE$ and the last two rows (except the biggest one in the fourth row) with permission from movie “The Man From Earth”.

the example shown in Figure 12, the subtitles generated that way may block some faces. In contrast, our subtitle placement procedure is more robust for complex backgrounds. The local energy term in Equation (5) tends to position the subtitle close to the speaker and at the same time far away from nonspeakers, thus lessening confusion and reducing eyestrain. Furthermore, our subtitle placement algorithm is consistent across frames thanks to the global energy term in Equation (6).




Fig. 11. Performance comparison in terms of speaker detection accuracy. Left: performance improvement of different features. Right: performance improvement compared with the previous method. Video indexes 1 to 8 refer to the TV clips from “Friends.S10E15” and “The.Big.Bang.Theory.S05E16”, and the movie clips from “Kramer.vs.Kramer”, “Erin.Brockovich”, “Up.In.The.Air”, “The.Man.From.Earth”, “Scent.of.a.Woman” and “Lions.for.Lambs” respectively. Ground truth speakers are manually labeled in our experiments.



Fig. 12. Use of saliency map in a complex background. Left: original screenshot; Middle: a saliency map used in the method of Hong et al. [2010]; Right: subtitle placement by our method (yellow shadowed rectangle) vs. that by Hong et al. [2010] (green shadowed rectangle) for the speaker (yellow rectangle), red rectangles are nonspeakers. The subtitle placement by the method of Hong et al. [2010] mistakes the nonspeaker #4 to be speaking and blocks nonspeaker #3, thus confusing the viewer. Screenshots with permission from YouTube video $v=tpBtdKRDVbI$.

4.3. User Study

To evaluate our system, we conducted a comprehensive usability study⁷ to compare it with conventional fixed position subtitles and a previous dynamic subtitling method.

4.3.1. Video Clips Selection. For the usability study, we selected 11 video clips from the above eight TV/movie videos. Specifically, two clips were selected from the movies *Erin Brockovich*, *Up in the Air*, and *The Man from Earth* and one clip from each of the other five videos. Video clips were selected from video segments containing several speakers with significant switching of dialog. The average length of the 11 video clips is 2.3 minutes.

For a fair comparison, the video clips were chosen such that the speaker detection accuracy of them is consistent with that of the entire TV/movie videos from which they were selected ($\pm 1.84\%$ in our user study). The consistency here is only considered for the video clips selected for our results (*Dynamic III*) instead of to all versions (*Static*, *Dynamic I* and *Dynamic II*) which was probably not possible. Video clips of other versions were selected accordingly.

4.3.2. Video Group Setup. For each of the 11 video clips, we produced a group of four different versions, namely *Static*, *Dynamic I*, *Dynamic II* and *Dynamic III*. Details of the version for each video group are shown in Table II.

Note that volume demonstration and subtitle highlighting are included in the method of Hong et al. [2010] because their method was mainly developed for people with hearing impairment. However, their usability study showed that the effects of volume demonstration and subtitle highlighting were minor for people with hearing impairments.

⁷https://sites.google.com/site/smartsubtitles/user_study.

Table II. Details of the Four Versions for Each Video Group

Video Versions	Description
Static	Traditional fixed version where subtitles are always put at the bottom of the screen
Dynamic_I	[Hong et al. 2010]
Dynamic_II	Combine our speaker detection with [Hong et al. 2010]'s subtitle placement
Dynamic_III	Ours




Fig. 13. Information about the participants in our user study (*gender, age, educational background and native language*).

Therefore, these cues were not used in the *Dynamic_I* version because it may be a distraction for the users in our study group. Users were asked to score (on a scale of 1 to 10) each of the four versions for overall viewing experience and for eyestrain level. The scores are uniformly defined with a higher score corresponding to a better overall viewing experience or a less eyestrain level.

4.3.3. Audio and Subtitle Language Setup. Based on the assumption that subtitles help the viewer understand the videos better, we chose specific videos with such audio so that subtitling would be useful rather than redundant. We chose clips with audio in French, Spanish and Hindi depending on the participant's language background. Subtitles, where possible, were in the native language of the reviewer, that is, Chinese for native Chinese speakers and English for native English speakers (or others).

4.3.4. Results. Each user was asked to view at least one video group of the four versions in entirety. In total, there were 219 participants in our usability study, amounting to about 20 for each video group. The users were of various ages and genders, from many different countries, and had different language and educational backgrounds (Figure 13).

The results of the user study are shown in Figure 14. We can see that all our results outperformed the *Dynamic_I* version [Hong et al. 2010] in terms of the overall viewing experience and degree of eyestrain level. In terms of overall viewing experience, the results showed that the *Dynamic_I* version [Hong et al. 2010] was worse or equal to the experience of Static version in 4 of the 11 groups. In comparison, the results from the *Dynamic_III* version (our method) outperformed the Static version in all groups. In terms of eyestrain, the method of Hong et al. [2010] was worse than or equal to the Static version in 2 of the 11 groups, whereas our method was better than the Static version in all but one group. Nevertheless, the results for that one group (video #9) were still better than the *Dynamic_I* version [Hong et al. 2010]. The *Dynamic_II* version outperformed *Dynamic_I* version in all groups for overall viewing experience and in 10 of the 11 groups for eyestrain. Our results indicate that our speaker detection is more accurate than the method of Hong et al. [2010]. Moreover, our results outperformed the *Dynamic_II* versions in terms of both overall viewing experience and eyestrain level, further validating our improved subtitle presentation method.

It should be noted that, in video #9 the level of eyestrain in all three dynamic subtitle versions were worse than in the Static version, although our method was the best. The

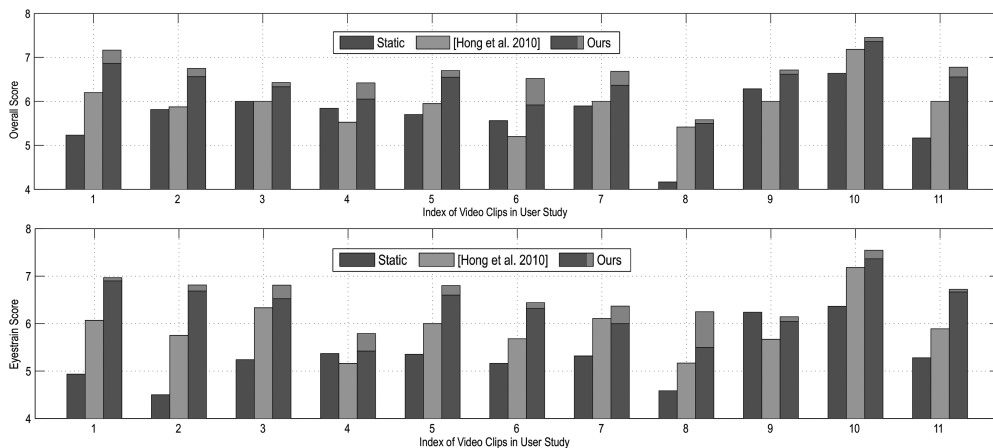


Fig. 14. Result of the user study. Overall viewing experience (above) and eyestrain level (below). Higher scores indicate better outcome. *Dynamic I* version is the method of Hong et al. [2010], the purple colored part of the third bar is for *Dynamic II* version and the whole third bar is our final version.

corresponding scene in the clip showed seven people arguing with the camera shot changing rapidly, which probably caused viewers to follow the rapid switching of the subtitle position in the changing conversation.

4.4. Applications

Our method can be used in several applications.

4.4.1. As a Tool to Assist the TV/Movie Industry. The usability study has demonstrated that our system can assist the TV/movie industry as a speaker detection tool for subtitling to provide a better viewing experience. Although not 100% accurate (>90% accuracy), our speaker detection tool can still replace more than 90% of the manual work in generating the speaker-following subtitles for TV/movies.

4.4.2. Better Presentation of Multimedia in Noisy Environments and for Users with Hearing Impairment. In noisy environments (e.g., subways and squares) or in a quiet public area, audio from the video may be hard to hear clearly or is muted. The viewer may then need speaker-following subtitles to know “who is speaking or what is being said” in order to better understand and enjoy the multimedia such as news, music videos, and television interviews.

As shown in the user study conducted by Hong et al. [2010], videos with dynamic subtitles can help users with hearing impairment understand the video contents. There are more than 66 million people suffering from hearing impairment, who can potentially benefit from the technology presented here.

4.4.3. Automatic Speech Balloon Layout for 2D/3D Game Avatars. Previous methods [Park et al. 2008a, 2008b] simply put speech balloons above game avatars’ heads. Our system can be used for automatic word balloon positioning for avatars in 2D/3D games to provide better presentation of speech for game avatars.

4.5. Limitations

Our system cannot always correctly handle scenes where there are no speakers such as subtitles from a gramophone or a telephone message. If the speaker detection algorithm does not detect a speaker, then our system can still correctly put the subtitle into the default bottom screen position. However, if people in the scene move in such a way

that confuses the speaker detection algorithm, our system may assign the subtitle to a wrong speaker, causing misunderstanding.

5. CONCLUSION AND FUTURE WORK

To enhance video viewing experience, we have developed an automatic system to place video subtitles next to their corresponding speaker. The proposed system detects speakers and then computes optimal positions to place the subtitles. Improving the performance of both the speaker detection step and subtitle placement step enables a better presentation of video subtitles in TV and movies. A comprehensive usability study with 219 users validated the effectiveness of our system. We have also suggested several possible applications of our system.

Although our system consistently outperforms both static subtitling and a previous dynamic method, there is still room for improvement to increase user satisfaction in terms of overall viewing experience and eyestrain level. In this regard, we will further improve speaker detection accuracy especially in more challenging situations (e.g., the TV series *ER*, in which surgeons and nurses wear face masks and move around all the time). We also aim to extend our system to cope with sound-to-text libraries in online applications such as teleconferencing and real-time video streaming where no subtitle information is available.

ACKNOWLEDGMENTS

The authors would like to thank all the participants of the usability study for their time and valuable feedback and comments.

REFERENCES

- X. Anguera Miro, S. Bozonnet, N. Evans, C. Fredouille, G. Friedland, and O. Vinyals. 2012. Speaker diarization: A review of recent research. *IEEE Trans. Audio Speech Lang. Process.* 20, 2, 356–370.
- B. Chun, D. Ryu, W. Hwang, and H. Cho 2006. An automated procedure for word balloon placement in cinema comics. *Adv. Visual Comput.* 576–585.
- A. Dimou, O. Nemethova, and M. Rupp 2005. Scene change detection for h. 264 using dynamic threshold techniques. In *Proceedings of the 5th EURASIP Conference on Speech and Image Processing, Multimedia Communications and Service*.
- J. Driver. 1996. Enhancement of selective listening by illusory mislocation of speech sounds due to lip-reading. *Nature* 381, 6577, 66–8.
- M. Everingham, J. Sivic, and A. Zisserman. 2006. Hello! my name is... buffy—automatic naming of characters in TV video. In *Proceedings of the 17th British Machine Vision Conference (BMVC'06)*.
- W. Faernando, C. Canagarajah, and D. Bull. 2001. Scene change detection algorithms for content-based video indexing and retrieval. *Electron. Commun. Eng. J.* 13, 3, 117–126.
- M. Gordan, C. Kotropoulos, and I. Pitas. 2002. A support vector machine-based dynamic network for visual speech recognition applications. *EURASIP J. Appl. Signal Process.* 1, 1248–1259.
- R. Hong, M. Wang, M. Xu, S. Yan, and T.-S. Chua. 2010. Dynamic captioning: video accessibility enhancement for hearing impairment. In *Proceedings of the International Conference on Multimedia*. ACM, 421–430.
- G. Jaffré, P. Joly, et al. 2004. Costume: A new feature for automatic video content indexing. In *Proceedings of RIAO: Coupling Approaches, Coupling Media and Coupling Languages for Information Retrieval*. 314–325.
- M. A. Just and P. A. Carpenter. 1987. *The Psychology of Reading and Language Comprehension*. ERIC.
- C. Kuo, C. Huang, and R. Nevatia. 2010. Multi-target tracking by on-line learned discriminative appearance models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 685–692.
- D. Kurlander, T. Skelly, and D. Salesin. 1996. Comic chat. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*. ACM, 225–236.
- G. W. McConkie, P. W. Kerr, M. D. Reddix, D. Zola, and A. M. Jacobs. 1989. Eye movement control during reading: II. Frequency of refixating a word. *Perception Psychophysics* 46, 3, 245–253.

- G. Monaci. 2011. Towards real-time audiovisual speaker localization. In *Proceedings of the 19th European Signal Processing Conference*.
- H. Nock, G. Iyengar, and C. Neti. 2003. Speaker localisation using audio-visual synchrony: An empirical study. In *Image and Video Retrieval*, 565–570.
- S.-H. Park, S.-H. Ji, D.-S. Ryu, and H.-G. Cho. 2008a. A smart and realistic chatting interface for gaming agents in 3-d virtual space. In *Proceedings of the International Conference on Games Research and Development*.
- S.-H. Park, S.-H. Ji, D.-S. Ryu, and H.-G. Cho. 2008b. A smart communication system for avatar agents in virtual environment. In *Proceedings of the International Conference on Convergence and Hybrid Information Technology (ICHIT'08)*. IEEE, 119–125.
- G. Potamianos, C. Neti, G. Gravier, A. Garg, and A. Senior. 2003. Recent advances in the automatic recognition of audiovisual speech. *Proc. IEEE* 91, 9, 1306–1326.
- K. Rayner. 1975. The perceptual span and peripheral cues in reading. *Cognitive Psych.* 7, 1, 65–81.
- K. Saenko, T. Darrell, and J. Glass. 2004. Articulatory features for robust visual speech recognition. In *Proceedings of the 6th International Conference on Multimodal Interfaces*. ACM, 152–158.
- K. Saenko, K. Livescu, M. Siracusa, K. Wilson, J. Glass, and T. Darrell. 2005. Visual speech recognition with loosely synchronized feature streams. In *Proceedings of the 10th IEEE International Conference on Computer Vision (ICCV'05)*. Vol. 2, IEEE, 1424–1431.
- I. K. Sethi and N. V. Patel. 1995. Statistical approach to scene change detection. In *Proceedings of IS&T/SPIE's Symposium on Electronic Imaging: Science & Technology*. International Society for Optics and Photonics, 329–338.
- M. Uricár, V. Franc, and V. Hlavác. 2012. Detector of facial landmarks learned by the structured output svm. In *Proceedings of the 7th International Conference on Computer Vision Theory and Applications*. Vol. 1, 547–556.
- P. Viola and M. Jones. 2004. Robust real-time face detection. *Int. J. Comput. Vision* 57, 2, 137–154.
- M. Wallace, G. Roberson, W. Hairston, B. Stein, J. Vaughan, and J. Schirillo. 2004. Unifying multisensory signals across time and space. *Exp. Brain Research* 158, 2, 252–258.
- Wikipedia. 2012. Vision span.
- L. Zhang, L. Chen, M. Li, and H. Zhang. 2003. Automated annotation of human faces in family albums. In *Proceedings of the 11th ACM international conference on Multimedia*. ACM, 355–358.

Received October 2013; revised February 2014; accepted April 2014