

Coding Instructions for Topic Segmentation of the AMI Meeting Corpus

(version 1.1)

Weiqun Xu, Jean Carletta, Jonathan Kilgour and Vasilis Karaiskos*

School of Informatics, University of Edinburgh
Email: {wxu,jeanc,jonathan,vkaraisk}@inf.ed.ac.uk

June 9, 2005

ABSTRACT

This paper is a manual that instructs annotator how to work on topic segmentation in the AMI meeting corpus. After introducing some general ideas of the task and the tool, it explains how to do the job step by step.

Contents

1	Introduction	1
2	The Task	2
2.1	Segmentation	2
2.1.1	Sub-topics	2
2.2	Topic Description	2
2.2.1	Standard “Topic” Descriptions	3
3	The coding tool: ICSI topic segmenter	4
3.1	Hints for finding segment boundaries	6
3.2	Other tips for coding	7

1 Introduction

The AMI Meeting Corpus is a collection of recordings of meetings based on a scenario. The subject of this scenario is the design of a remote control device in a virtual company, Real Reactions. The task is carried out in a series of four meetings by four participants that take the roles of a project manager, a marketing expert, a product designer, and a user interface designer. The participants are given information about their roles, the task in question and what happens between two subsequent meetings.

Many people study these meetings because one of the big technology challenges at the moment is to build a meeting browser — that is, something that can be used to find out what happened at a meeting. A big part of meeting browsing is knowing what the people in a meeting were talking

*Contact person. All kinds of comment and feedback are warmly welcome.

about — the topic — and when they changed topics. It’s easier to make a machine understand topics and topic changes if someone tells it about topics and topic changes on some examples. Your job is to listen to some meeting recordings, divide the meetings up by topic into “segments”, and briefly describe what the topic is for each one. We expect this to take 4-5 hours for each meeting.

2 The Task

2.1 Segmentation

You might be expecting us to tell you exactly how many topic segments each meeting should have, but the truth is, it varies considerably. We have in mind that a typical one-hour meeting might have something in the order of six to ten segments, but there could be meetings that discuss one topic extensively, and others that handle a very large number of topics, all briefly. There is also no optimal length for a topic; there is a fine, yet subjective balance as to how many topics one could detect in a single meeting, before it all seems too fragmented. For this reason, you should divide the meeting into segments in the way that you find most natural. By this we mean, that short-lived deviations off a main point of discussion, for example, do not have to be marked as a separate topic. Depending on the nature and, even more, on the length of the deviation, it can form a segment of its own, or ignored and incorporated in the surrounding segment.

Everything that is said during a meeting should end up in some segment, and since the meetings are based on a given scenario, we expect that — at least most — topics do recur. For this reason we provide some standard kinds of segments (see 2.2.1).

2.1.1 Sub-topics

Sometimes you won’t be sure whether to mark part of a meeting as one segment or two, because there are really two segments but they are related to each other. For instance, if a group were talking about what they liked about Edinburgh, they might talk first about the free museums and then they might talk about the architecture. If they talked about these two things completely separately, without saying anything about why the two go together, then they would be separate topics. However, if they made clear that there were some connection, for instance, by saying that they were talking about what they liked about Edinburgh and then introducing these themes, the overall segment would be about “good things about Edinburgh”, with the sub-segments about “free museums” and “architecture”. You can mark sub-segments wherever you like, to cover part of the material in a segment.

If you feel that there is a clear subtopic happening, then mark it and describe it. Otherwise, don’t bother to subdivide. In theory every time someone talks they’re saying something different from the last person and therefore it should be possible to mark a new subtopic, but, as we mentioned previously, we don’t want this level of detail. A sub-segment should be something that the group is discussing, not just something one person threw into the discussion. We would expect some sub-topics in the meeting corpus, and possibly (but very rarely) some sub-sub-topics; if you find highly nested structures (with lots of detailed sub-sub-topics), you should consider whether you might be subdividing topics too finely.

Finally, if you decide that a part of a topic A should form a sub-topic, you should not feel obliged to assign the rest of A to other sub-topics; they can simply be subsumed to A only (see also section 3).

2.2 Topic Description

As well as saying where the discussion of a topic starts and ends, you need to give a short description of the topic. This can usually be based on a few keywords from the discussion, and needs to be detailed enough that someone could figure out later on what the topic was. There already is a list of topic descriptions (see Table 1) to choose from. Once you segment a topic, you should first check this

list and see whether you can find a fitting description. Only if this list does not contain a label that seems appropriate, should you create a new one.

Groups will quite often discuss a topic and then return to it later in the same meeting. When this happens, you should use exactly the same description. The software lets you do this without typing it in again. We'll know which topic a subtopic goes with, so it's OK if it is necessary to read both the topic and subtopic label to understand what the subtopic is (e.g., "why we like Doris" could have "her hairstyle" for a subtopic rather than "hairstyle as a reason for liking Doris").

2.2.1 Standard "Topic" Descriptions

The scenario defines the structure of each meeting. In the first meeting of each set, after the introductions and explanations about the project goals, the participants spend some time drawing animals on the whiteboard, in order to get acquainted with the equipment. In the second and third meeting, there is a presentation by each participant (except for the project manager), followed by a discussion and decisions on whatever was suggested. The second meeting also includes the announcement of some additional design requirements and the definition of the user group the product will be targeted at. The fourth meeting includes a presentation and evaluation of one or more prototypes created, estimating the costs and budgets and a discussion on what the participants think about the whole desing process.

We provide a pre-set list of topic descriptions to make coding easier and to ensure some level of consistency in the descriptions across the coded meetings (Table 1). We have divided the descriptions into three categories:

TOP-LEVEL TOPICS refers to topics whose content largely reflects the meeting structures as described above; it is expected that you will found these in every set of meetings you annotate; it is unlikely that during the course of the annotation, this list will need to be expanded;

SUB-TOPICS are what parts of the TOP-LEVEL TOPICS may form; the labels you have here are quite general; however, you may find during the course of the annotation that this list needs to be populated with more descriptions;

FUNCTIONAL descriptions, on the other hand, generally refer to these parts that either refer to the very process and flow of the meeting, or are simply irrelevant; they can be either a top-level topic or a sub-topic, but you shouldn't try to identify any sub-topics within them.

For every topic segment that you need to label, you can use any label from these lists, or, if none of them fit well, you can make up your own new label. On the other hand, just because a certain description is in the list, it does not mean that you must necessarily identify such a topic in any given meeting. In any case, the main rule is that if, when you select a topic, one of these labels seems like a good description for it, then use it; otherwise, you can always create a new description.

So, descriptions for any top-level topics you segment are expected to be found either in **TOP-LEVEL TOPICS** or **FUNCTIONAL** lists. On the other hand, any sub-topics will get their labels either from **SUB-TOPICS** or **FUNCTIONAL**.

However, the **FUNCTIONAL** labels require some further clarification:

Opening

The "opening" topic description can be used for all of the little things that groups tend to do at the beginning of a meeting: take attendance, review the minutes of the previous meeting or set the agenda, say how long the meeting is, and so on. It's useful for us to know where the opening begins and ends, but the opening doesn't contain particular topics, or at least not technical ones, just discussion about how to conduct the meeting. We don't need opening material to be further segmented.

Closing

Just as meetings can have openings that aren't about a real topic, they can also be closed in a similar way, for instance, with time dedicated to setting the next meeting date, briefly recapping and

TOP-LEVEL TOPICS	SUB-TOPICS	FUNCTIONAL
project specs and roles of participants*	project budget	opening
new requirements**	existing products	closing
user target group**	trend watching	agenda/equipment issues
interface specialist presentation***	user requirements	chitchat
marketing expert presentation***	components, materials and energy sources	
industrial designer presentation***	look and usability	
presentation of prototype(s)****	how to find when misplaced	
discussion***		
evaluation of prototype(s)****		
evaluation of project process****		
costing****		
drawing animals*		

- * Found by default in **a** meetings.
- ** Found by default in **b** meetings
- *** Found by default in **b, c** meetings
- **** Found by default in **d** meetings

Table 1: List of suggested topic description for scenario meetings.

reviewing what was decided and who will do what in preparation for the next meeting. Mark these as “closing”, and again, do not segment further.

Agenda/Equipment issues

Throughout the meeting the discussion can be diverted towards agenda issues, like the order in which presentations, or what’s next in the list of decisions they have to take in the process of the meeting, etc.; similarly, the flow of conversation can also be interrupted by trying to work out the computer or some other piece of equipment. These parts of the conversation do not contain “useful” information. Therefore, when you think that such a discussion becomes substantial (this is subjective and entirely up to your intuitions), segment the relevant part and mark it as “agenda/equipment issues”.

Chitchat

Sometimes during a meeting the participants just chat aimlessly, usually about social matters. This especially happens after the microphones have been switched on but before the beginning of the meeting “proper”, and again at the end. It can also happen in the middle of a meeting, for instance, when a projector breaks, or simply if someone drags the group off-topic. When this happens, divide the meeting so that these areas form their own segments, but label them with the special topic description, “chitchat”.

It can also happen that while the group is having a proper discussion of some topic, there will be one or two quick utterances, like jokes, that you might be tempted to code as “chitchat”, but the group doesn’t really get pulled off the topic they are discussing. Don’t bother to segment around these cases — just leave the utterances within the wider segment.

As mentioned above, there is no need to sub-divide any such **FUNCTIONAL** topics.

3 The coding tool: ICSI topic segmenter

We’ve written software specifically for this task that will allow you to view a transcription, play the meeting recording, and segment the meeting and add the topic annotation. The coding tool will work on uncoded, coded, or partly coded meetings, so you can stop and restart at any time (but remember

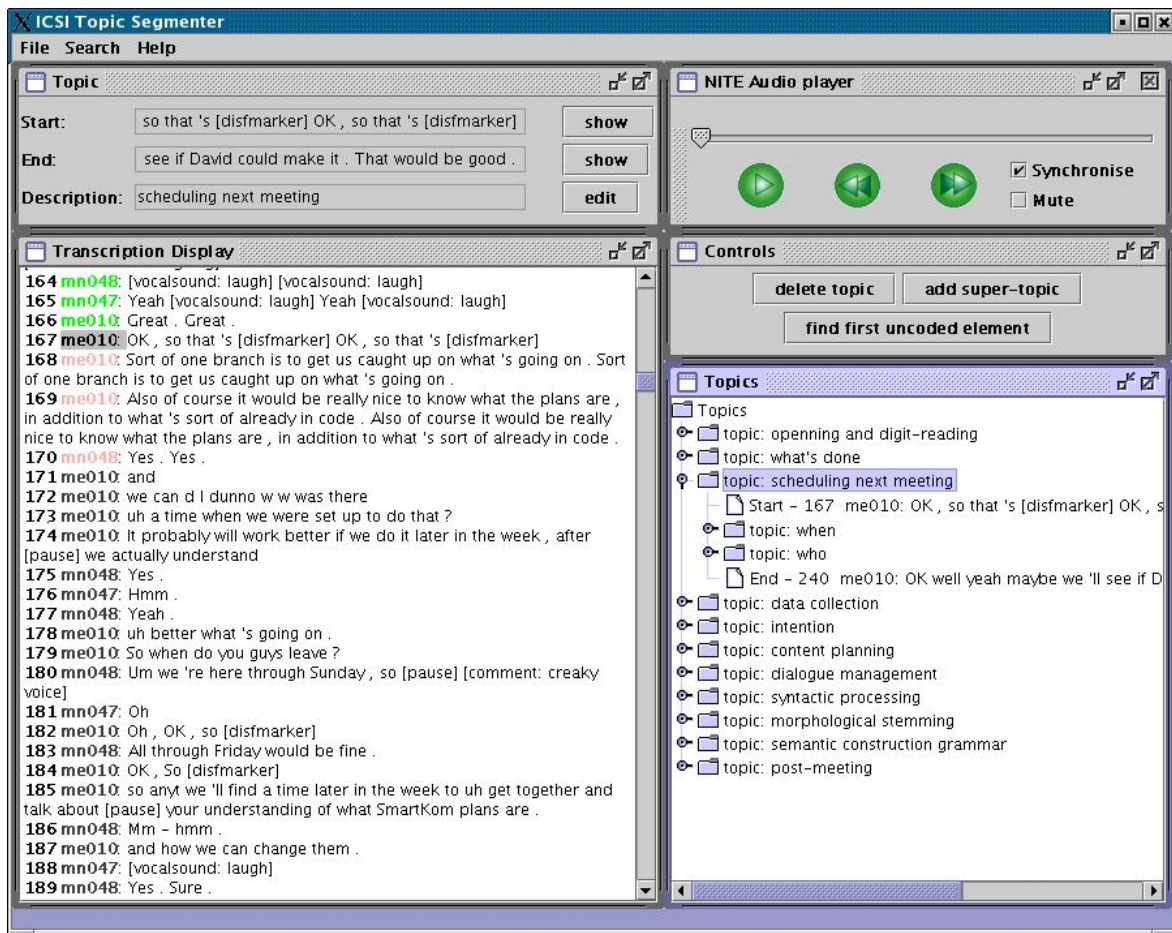


Figure 1: Snapshot of the coding tool

to save your work!) or just review meetings you coded earlier.

This section assumes that the tool has been installed and the data set up on your machine. A different document also provides the necessary information on how to load a particular piece of data on the tool ¹. Here, we shall explain how to use the different functions of the topic segmenter.

Opening a meeting usually results in five windows on a common desktop (see Figure 1):

TOPIC: which shows some information about the current topic segment, including its START and END and its topic DESCRIPTION. You may quickly move to the corresponding utterances by clicking the SHOW button. You can also change the description by clicking the EDIT button. When you click the EDIT button, a DESCRIBE TOPIC window will pop up, in which you can either *Choose an existing topic* (predefined or previously coded) from a drop down list or *add a new one* by inputting some free description.

NITE AUDIO PLAYER: which plays the audio of the opened meeting ². There are three buttons, two check boxes, and one progress bar.

¹For installation and data set-up issues, please contact your manager; the necessary information is at <http://wiki.idiap.ch/ami/IcsiInstallationInstructions>. All other information on using the annotation tool, as well as the complete annotation procedure is found at <http://wiki.idiap.ch/ami/TopicSegmentation>

²If there are problems with the sound card or the coding tool can not locate the corresponding audio file, the player may not work properly or even not appear. Contact your manager if this happens.

Progress bar – which shows audio play progress. You may also slide forward or backward to your desired part.

Buttons ▷ = play or || = pause, ▷▷ = fast forward, ◀◀ = fast rewind.

Check boxes

Synchronise when checked (default), audio is synchronised with text highlighted in green. This is very helpful when you browse the meeting for the first time, but might become annoying when you just want to skim the text later. This feature can be disabled by unchecking the box.

Mute turn off the sound while it’s playing.

TRANSCRIPTION DISPLAY: which displays the transcription of the opened meeting. Every utterance is preceded by a speaker ID label. The transcription window divides up the meeting by who said what. Beyond that, the way it divides something one person said into lines is fairly meaningless. Topic boundaries can be placed on any word in the transcript; an utterance can belong wholly to one segment or be split between two segments, if you judge that a topic shift happens within it.

If you want to add a topic segment left-click on first word of the segment: A “||” sign will appear to mark the beginning of the topic. Then right-click on the last word in the segment, which has to be after the first word. When there is a valid right click, a “;” sign will appear at the end of the word marking the end of the topic. The interface will also pop up a window for you to describe its topic. You can do that either by choosing one from the drop-down list (see Table 1) or by entering a new description.

If you want to start play the audio from an arbitrary line, select the line, then press **Ctrl** and right click at the same time.

CONTROLS³ : which provides additional ways of changing the topic coding and moving around the transcript.

Delete Topic: which deletes the topic (i.e., a segment and its topic description) after you click a topic in the TOPICS window. (Note: any sub-structure will not be deleted, instead all the children of the deleted topic will be upgraded one level.) The interface is slow to adjust after you delete a topic, so be patient.

Add Super-Topic: which adds a super/parent topic. First, select multiple consecutive topics in the TOPICS window. (You can do this by selecting the first topic in the series with the left mouse button, holding down the *shift* key, and selecting the last topic in the series.) Then click this button to make the selected topics sub-topics of some new parent topic. Of course, you need to give a description of the super-topic.

Find First Uncoded Element: which helps you quickly jump to the first uncoded utterance.

TOPICS: which displays all the coded topics in the meeting. A topic node is represented by its description with its start and end utterances and possible sub-topic(s). You can do nothing in this window but select one or more topic nodes.

3.1 Hints for finding segment boundaries

There are some clues that should help you find segment boundaries.

The first is that people in meetings quite often announce topic changes. For example,

Okay, now we’re moving on to [...] Marketing. [ES2008c]

You should be careful that the meeting actually moves on to the next topic at this point — sometimes someone will intervene with more material on the previous topic — but otherwise such utterances are pretty good indicators of a boundary. Note that as well as signalling a topic shift, these utterances often give some clear idea about the topic which you can put into topic description.

The second clue is that if the group discusses the agenda at the beginning, it can be useful to look for where the agenda items appear, even if groups don't always follow the agenda that they set. Again, the agenda can be a useful source of topic descriptions.

The last set of clues are words like “anyway” and “so”, which can be used to indicate a topic shift, like in the example above. When you listen to the recordings, these indicators can sound quite distinctive from other ways of using the same words.

3.2 Other tips for coding

When you start on a new meeting, we suggest that you begin by listening, using the transcription to skip forward once you have a sense of what's going on. If you're in a room with other people please use headphones. You can code any part of the meeting at any time, but it's best to work in an orderly fashion, from beginning to end. If subtopics feature heavily in the meeting you're coding, you may find it easier to segment the entire meeting and then return to add the subtopics afterwards.

When you finish or if you have to interrupt the segmentation, do remember to save your work (File/Save menu). Make sure that every word in the transcript has been assigned to a top-level topic and a description has been given for all topics. If you try to exit the program, you may get the following message:

```
You have not completed coding topics for this observation. Are you sure you want to stop?
```

This message appears both when there are parts of the meeting that have not been assigned to a top-level segment and when you have subdivided a top-level topic, but not every word in it has been assigned to a sub-topic. In the former case, you will have to return and complete the task. In the latter case, you can ignore the message, since you do not have to fully subdivide a topic (see also 2.1.1).

Finally, please remember to keep track of your hours and update any progress tables that have been set ⁴.

⁴All procedural information regarding this task can be found at <http://wiki.idiap.ch/ami/TopicSegmentation>.