Social Signal Detection in Spontaneous Dialogue Using Bidirectional LSTM-CTC

Hirofumi Inaguma[1] Koji Inoue[1] Masato Mimura[1] Tatsuya Kawahara[1]
[1]Graduate School of Informatics Kyoto University, Japan

Introduction

Goals

- To detect social signals robustly on the event-level rather than the frame-level

What's social signals ?

Speech cues (this study)
- ◆ Laughter
- ◆ Filler
- ◆ Backchannel
- ◆ Disfluency

Visual cues
- ◆ Facial expressions
- ◆ Gestures
- ◆ Postures
- ◆ Gaze

Social signal detection [Schuller+ '13]

✓ Useful for understanding speakers
✓ Informative for dialog systems to behave like human
✓ Rich annotation

Related works: frame-wise classifiers

☹ Does not directly lead to the event-unit detection [Gosztolya+ '15]
☹ Frame-level target labels are required (|inputs| = |outputs|)
☹ Post-processing are required (threshold or HMM etc.)

➡ CTC can solve all these problems!

Approach

Bidirectional Long-Short Term Memory (BLSTM)

✓ Aim for the accurate detection by using the future context

Connectionist Temporal Classification (CTC) [Graves+ '06]

✓ A loss function which can optimize sequence labeling where the input and the target label sequence have different lengths
✓ Works together with RNNs
✓ Removes the need to conduct segmentation
✓ Has potential of improving robustness of detection (spike prediction)

Key idea of CTC

1. Introduction of a *blank* label (−) (the network emits no labels)
2. Allow repetitions of the same labels

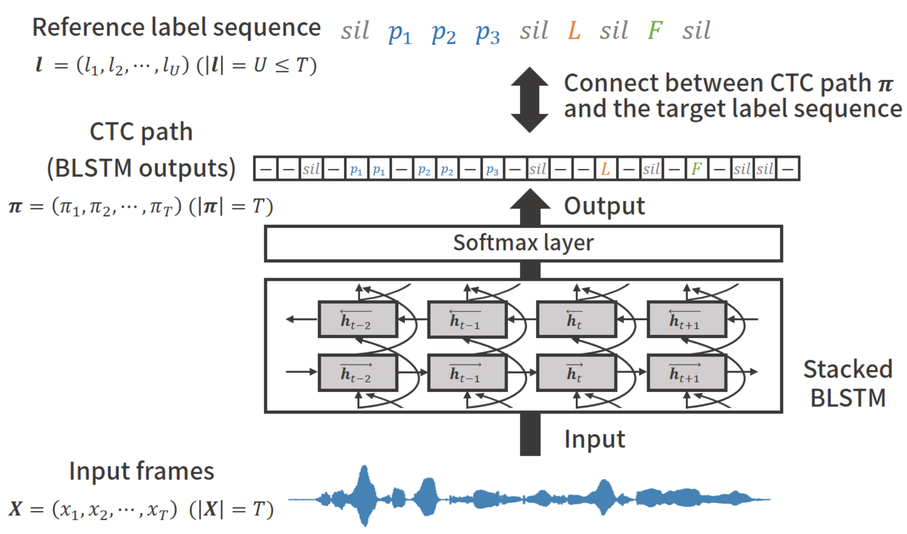

CTC Training

✓ Minimize $L_{CTC} = -\ln p(\boldsymbol{l}|\boldsymbol{X})$
✓ Marginalize $p(\boldsymbol{l}|\boldsymbol{X})$ by a summation of probability distribution of all possible frame-level alignments

$$p(\boldsymbol{l}|\boldsymbol{X}) = \sum_{\boldsymbol{\pi} \in \Phi^{-1}(\boldsymbol{l})} p(\boldsymbol{\pi}|\boldsymbol{X}) = \sum_{\boldsymbol{\pi} \in \Phi^{-1}(\boldsymbol{l})} \prod_{t=1}^{T} y^{t}_{\pi_t}$$

✓ Decompose $p(\boldsymbol{\pi}|\boldsymbol{X})$ based on the conditional independence assumption
✓ Compute $p(\boldsymbol{l}|\boldsymbol{X})$ efficiently with the forward-backward algorithm

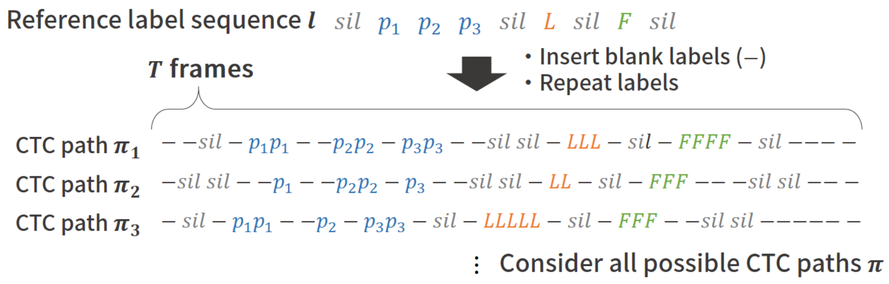

CTC Decoding

1. Remove repetitions
2. Remove all blank labels

Experiments

The SSPNet Vocalization Corpus (SVC)

✓ Used in Interspeech 2013 ComParE (total 8.4h) [Schuller+ '13]
✓ Laughter , Filler , Garbage (speech and silence)
✓ Not transcripts available in SVC
✓ Target labels corresponding to acoustic events in the input are required (speech or silence)

◆ Generation of training labels for CTC

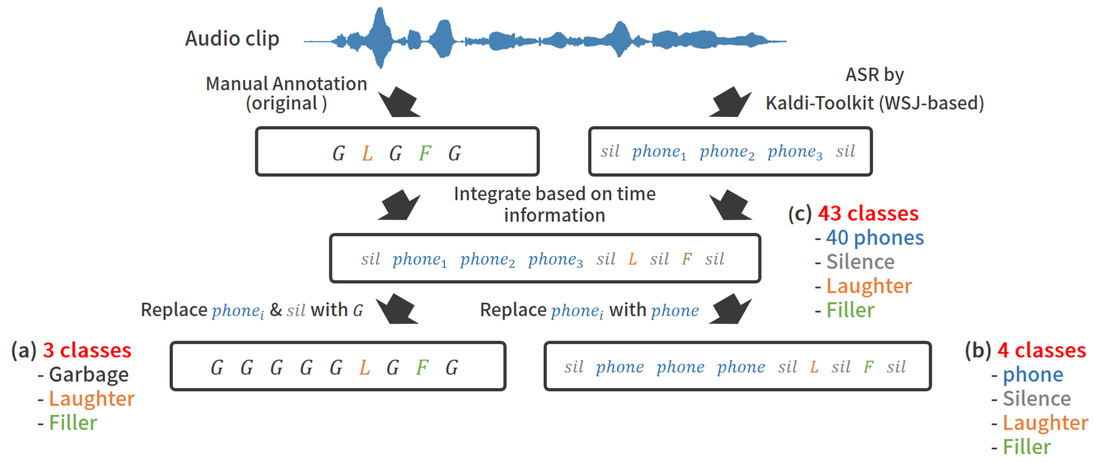

◆ Results

Class	Model	Laughter			Filler			Ave.
		Prec.	Rec.	F_1	Prec.	Rec.	F_1	F_1
3	AdaBoost-HMM [Gosztolya+ '15]	0.58	0.74	0.65	0.65	0.71	0.68	0.66
	DNN-HMM	0.58	0.72	0.64	0.71	0.60	0.65	0.65
	(a) BLSTM-CTC	**0.65**	**0.66**	**0.66**	**0.66**	**0.80**	**0.72**	**0.69**
4	(b) BLSTM-CTC	0.60	0.49	0.54	0.59	0.78	0.67	0.61
43	(c) BLSTM-CTC	0.79	0.51	0.62	0.71	0.78	0.74	0.68

CTC outperformed the conventional frame-wise classifiers even without time information in the training stage

◆ CTC outputs (posteriors)

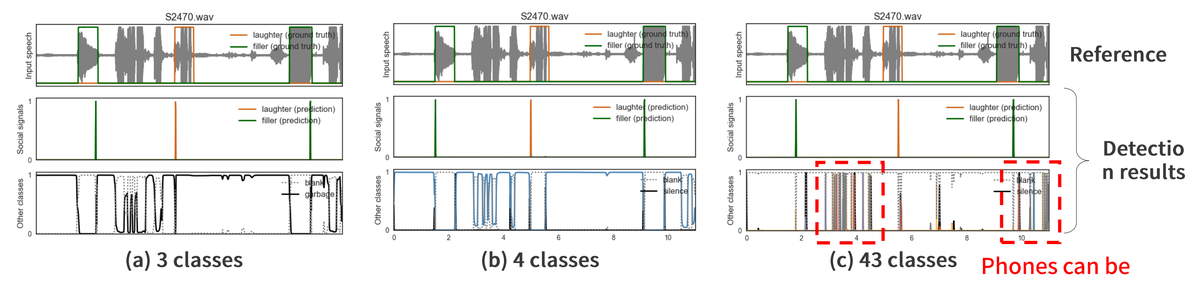

ERATO Human-Robot Interaction Corpus

✓ Japanese face-to-face spontaneous dialog with an android ERICA, which was remotely operated
✓ 91 sessions (about 10 min/session, total 16.8h)
✓ Laughter, Filler, Backchannel, Disfluency
✓ 4 social signals + 83 Japanese kana characters + space

◆ Generation of training labels for CTC

✓ Insert each social signal label in front of the corresponding word

$word_1$ ($Laughing$ $word_2$) $word_3$ → $word_1$ L $word_2$ $word_3$
$word_1$ ($Filler$ $word_2$) $word_3$ → $word_1$ F $word_2$ $word_3$
$word_1$ ($Backchannel$ $word_2$) $word_3$ → $word_1$ B $word_2$ $word_3$
$word_1$ ($Disfluecy$ $word_2$) $word_3$ → $word_1$ D $word_2$ $word_3$

◆ Results

Social signals	Prec.	Rec.	F_1
Laughter	0.89	0.35	0.50
Filler	0.75	0.75	0.75
Backchannel	0.86	0.87	0.86
Disfluency	0.44	0.15	0.22

Social signals	CER (%)
BLSTM-CTC (w/o social signals)	19.1
BLSTM-CTC (w/ social signals)	**18.6**

Joint-training with social signals improved character-level speech recognition accuracy

◆ CTC outputs (posteriors)

CTC could capture relationships between social signals and subwords

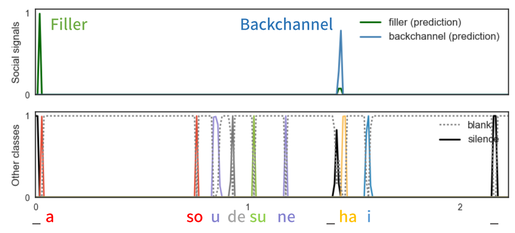

Conclusions

Summary

✓ Robust social signal detection on the event-level by BLSTM-CTC
✓ Removed the need of pre-alignment and post-processing
✓ Outperformed the conventional frame-wise classifiers
✓ Alignments are generally matched with the actual timing of the occurrence of social signal events

Future work

✓ Evaluate with large dataset
✓ Attention-based detection