Microsoft

Spatial audio feature discovery with convolutional neural networks

Etienne Thuillier, Hannes Gamper, Ivan J. Tashev
Microsoft Research Redmond

Need

Better understanding of elevation cues

... for efficient & immersive spatial audio rendering

Problem

Listening test studies

... *do not scale to large sample sizes*

... *are exposed to human error*

... *provide highly compressed data*

Proposed approach (2 steps)

1) Train a CNN to classify spatialized sounds

Dataset	Year	Subjects	Pairs/subj.
Aachen	2016	46	2304
Microsoft	2015	252	400
RIEC	2014	105	865
ARI	2010	135	1150
CIPIC	2001	45	1250

sample HRIR pair

spatialise 50 ms white noise burst

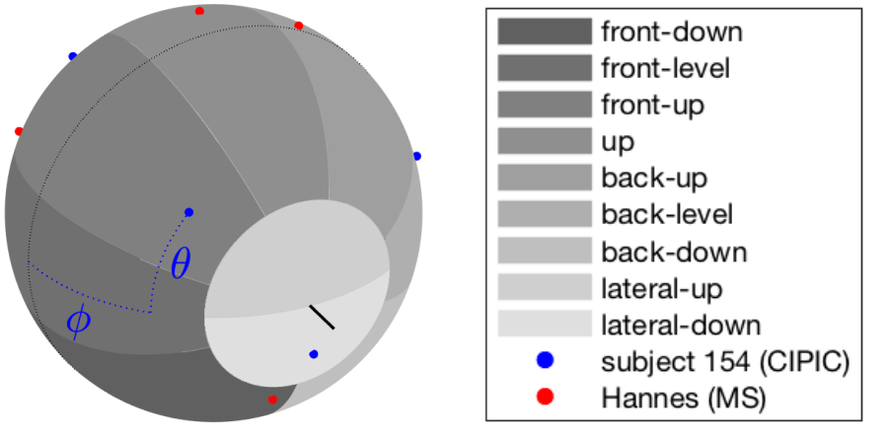

~500k HRIR pairs from 583 subjects!

2) Apply an explanation technique

[1]

GoogleNet

$$\begin{bmatrix} 0.05 \\ 0.10 \\ 0.75 \\ 0.10 \end{bmatrix}$$

'*sheep*'
'*halibut*'
'*shark*'
'*engineer*'

Deep Taylor Decomposition [1]

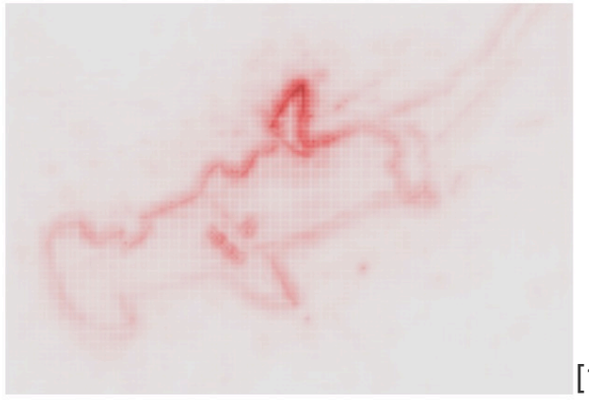

[1]

[1] G. Montavon et al., “Explaining nonlinear classification decisions with deep taylor decomposition,” *Pattern Recognition*, vol. 65, pp. 211–222, 2017.

2) Apply an explanation technique

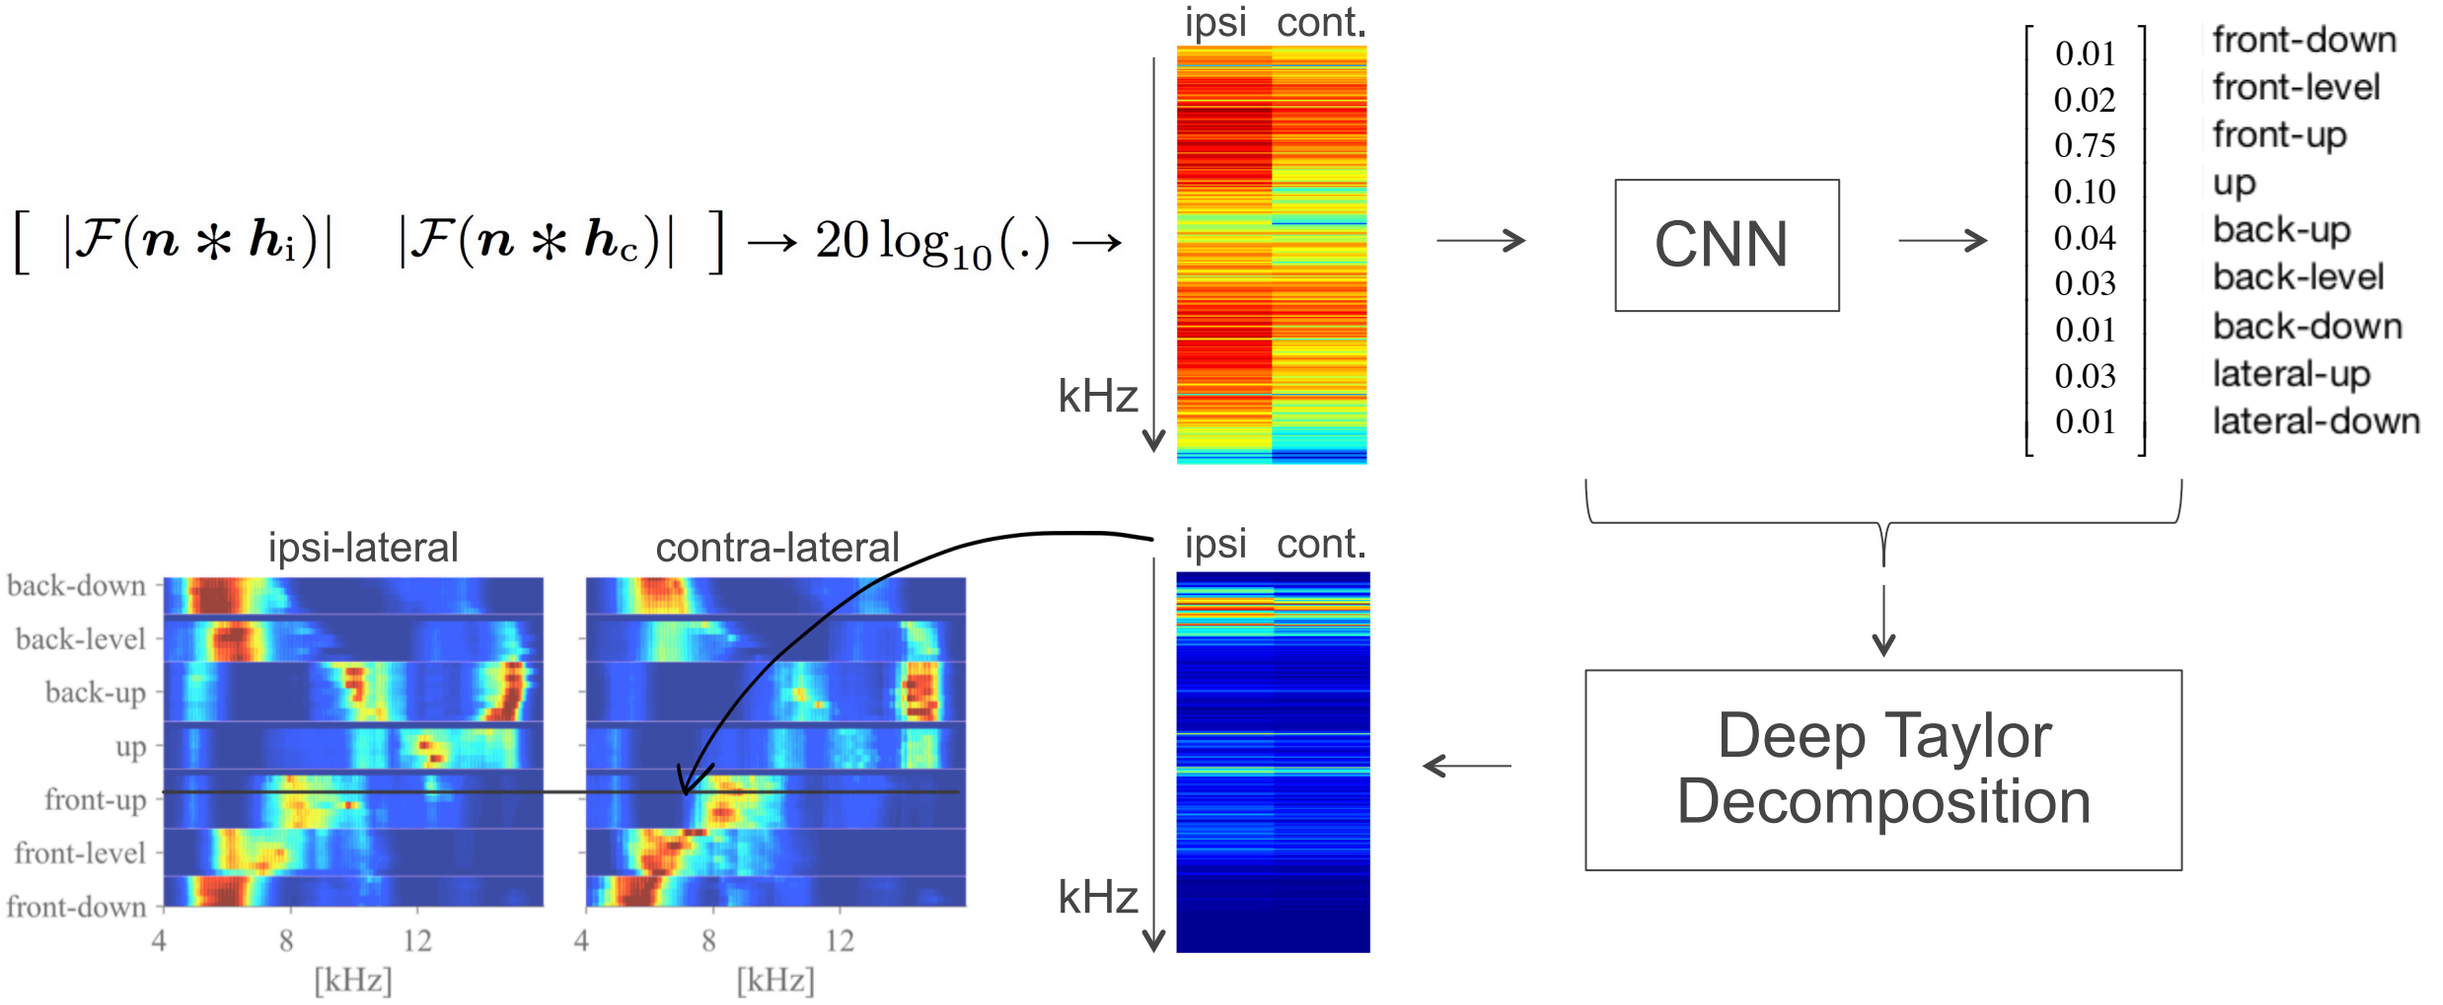

Models

Models

		input	**conv. 1**	conv. 2	conv. 3	**conv. 4**	filters/layer	stride	dense	output layer
WB	**> 300 Hz**	1005×2	**25×2**	11×1	11×1	**10×1**	4	2	216×9	soft-max
HP	**> 4 kHz**	1005×2	**25×1**	11×1	11×1	**10×2**	4	2	216×9	soft-max

< 3k trained parameters

> 500 training points parameter

~ 500 subjects in train set

~ 30% classification error (test)

	CE [%]	RMSE [deg]	MAE [deg]	r
random	91.3	74.5	59.5	0.65
[15]	-	25.2	-	0.85
[27]	-	-	22.3	0.82
[28]	-	-	≈25	-
WB	45.1	43.2	16.5	0.90

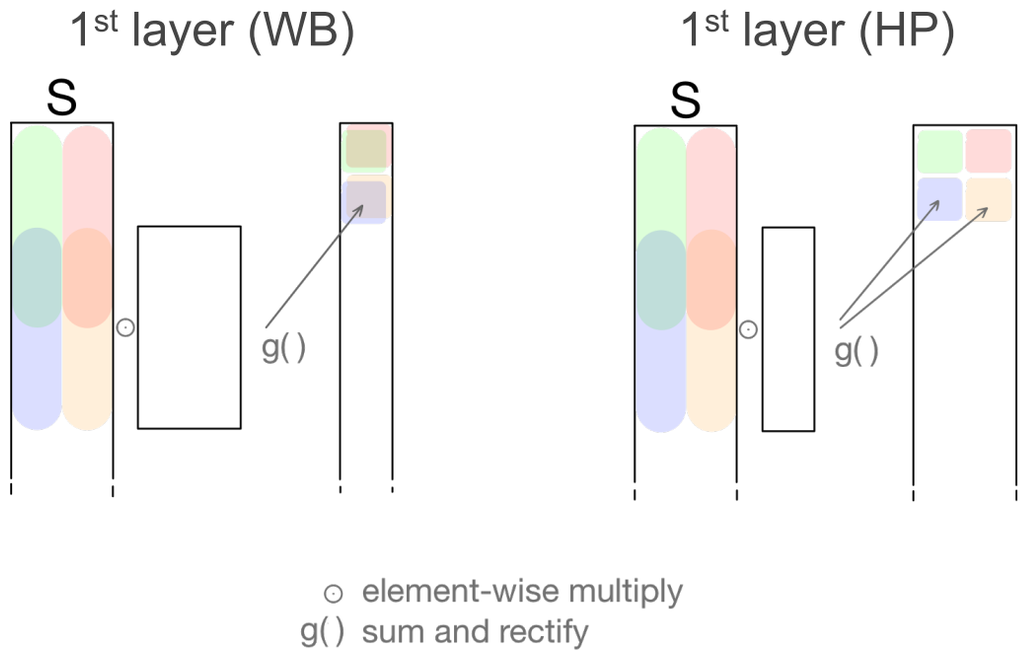

Results

HP model

Cone of confusion at 30º lat.

CIPIC's subject 154

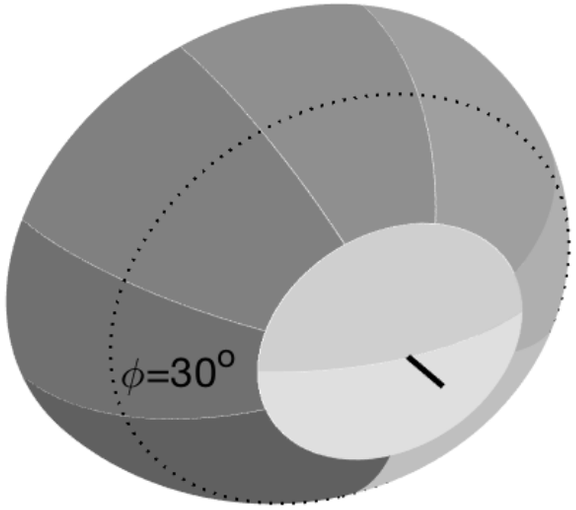

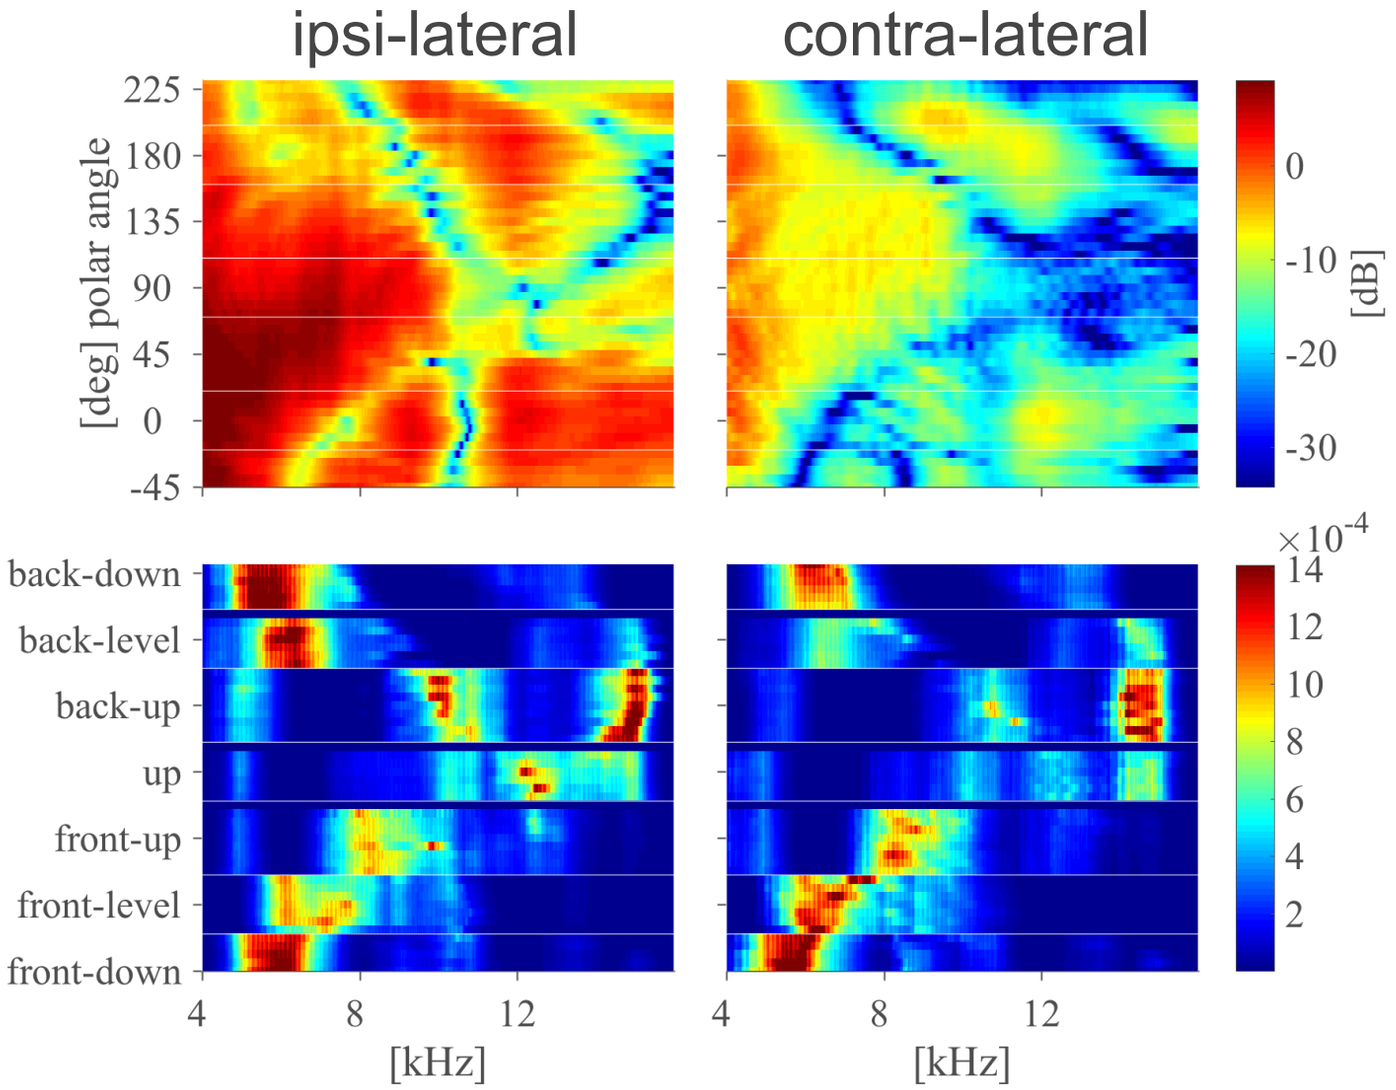

WB model

Cone of confusion at 30^{o} lat.

CIPIC's subject 154

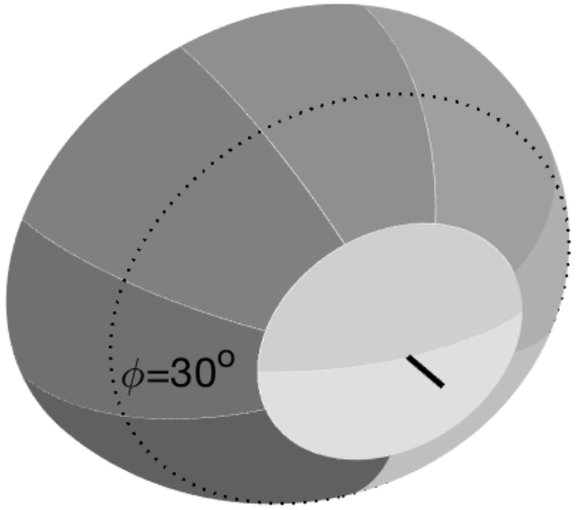

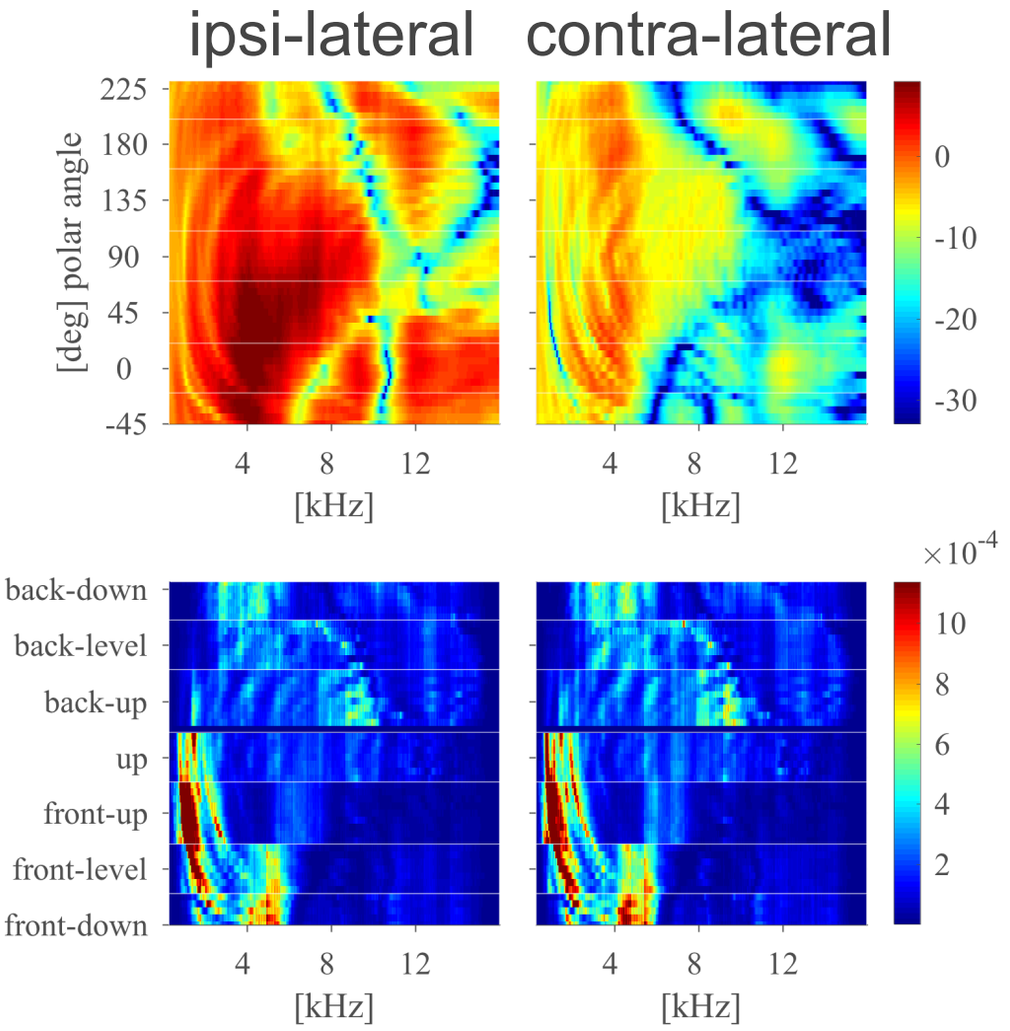

WB model

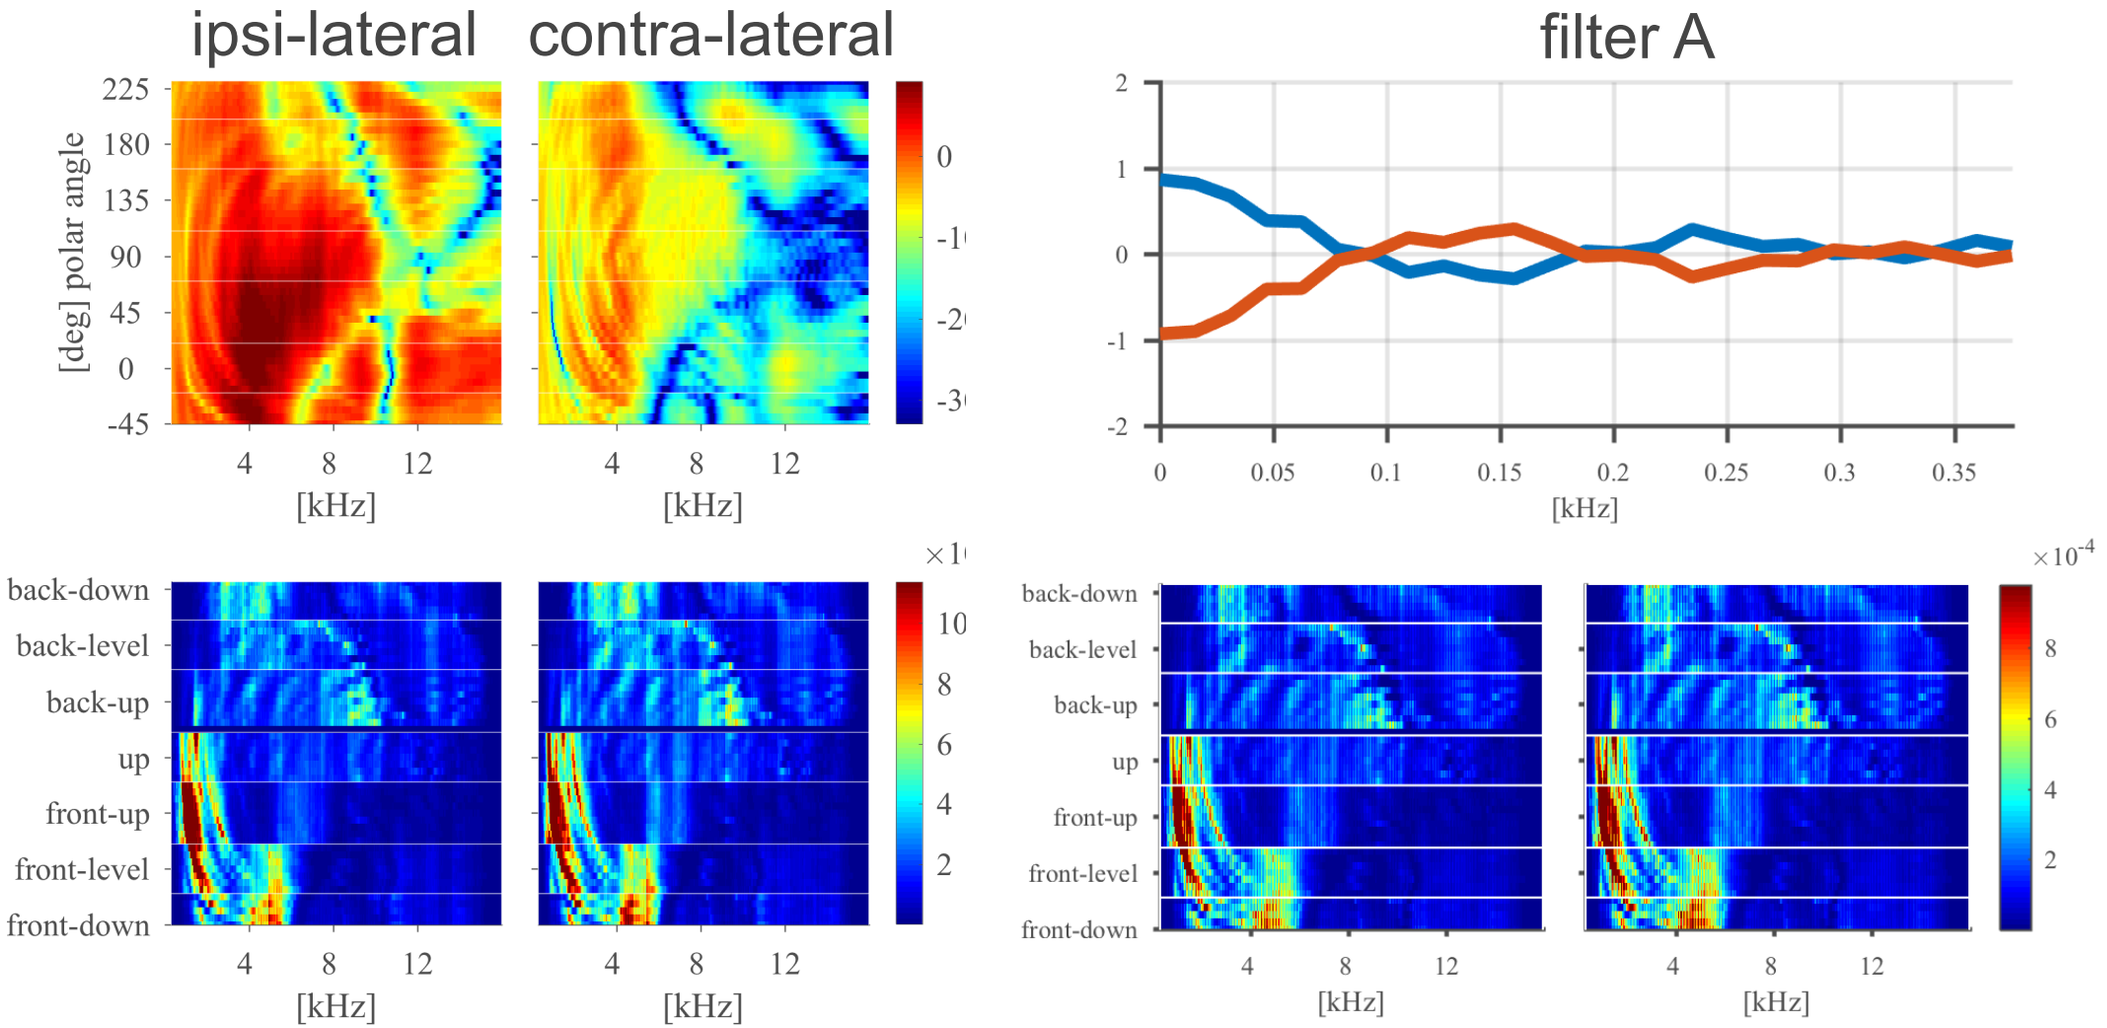
ipsi-lateral
contra-lateral
filter A
[deg] polar angle
225
180
135
90
45
0
-45
4
8
12
[kHz]
back-down
back-level
back-up
up
front-up
front-level
front-down
2
1
0
-1
-2
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
×10^-4

WB model

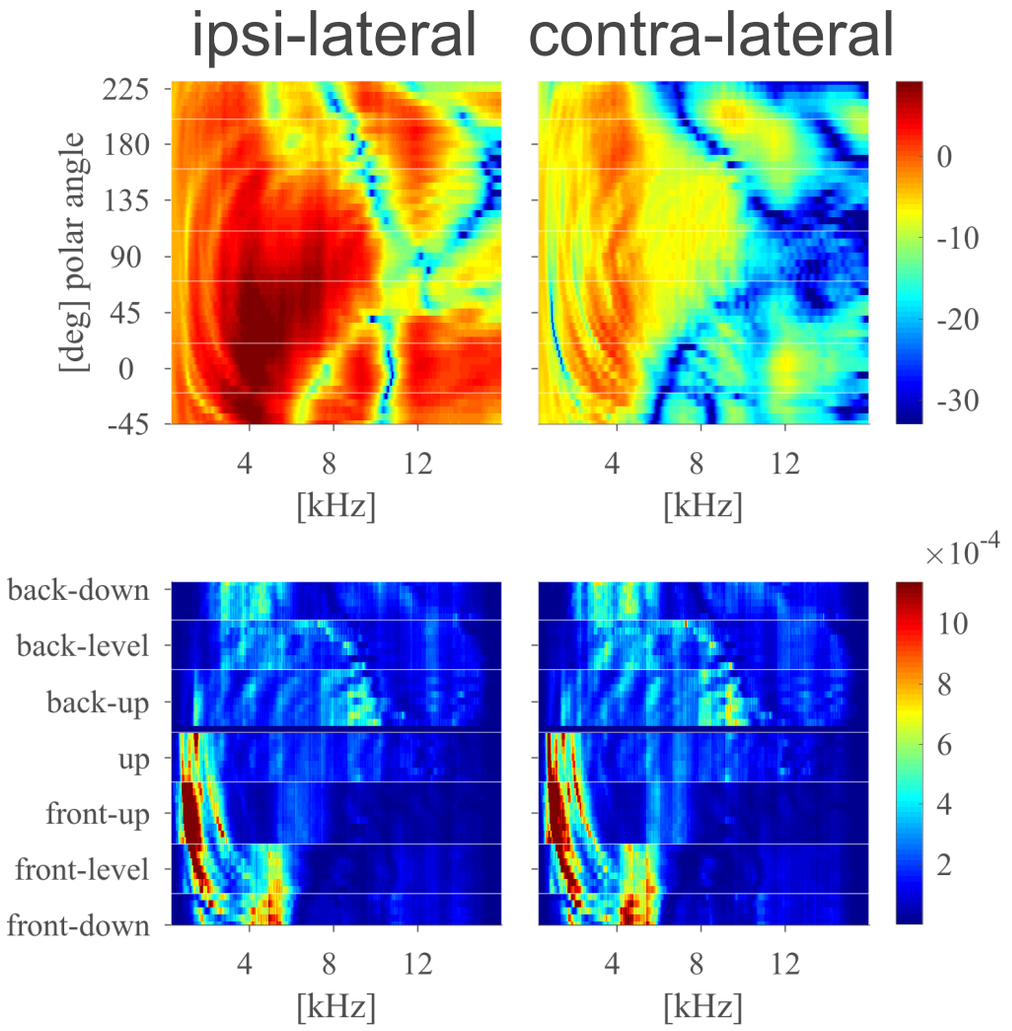
ipsi-lateral
contra-lateral
[deg] polar angle
225
180
135
90
45
0
-45
4
8
12
[kHz]
0
-10
-20
-30
back-down
back-level
back-up
up
front-up
front-level
front-down
×10-4
10
8
6
4
2

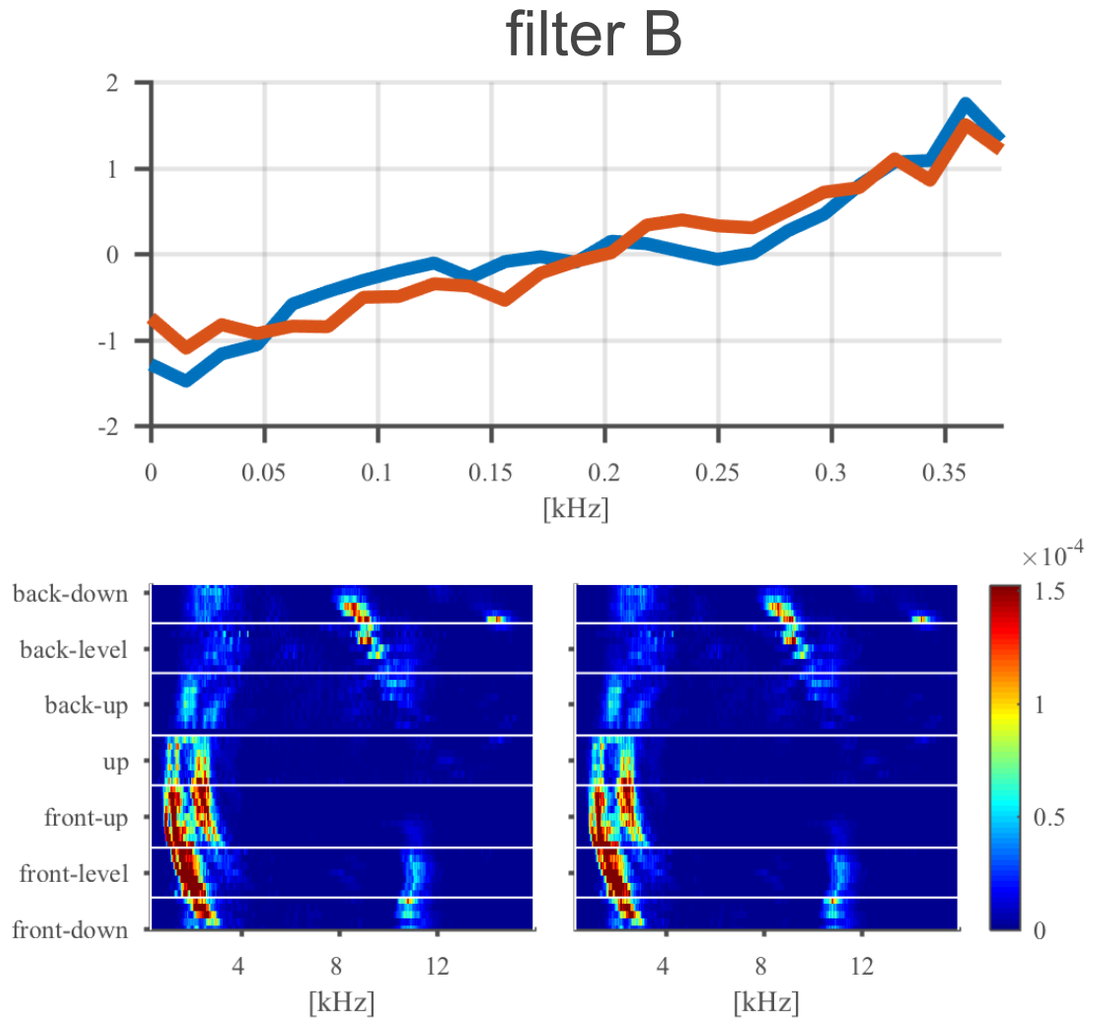
filter B
2
1
0
-1
-2
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
[kHz]
back-down
back-level
back-up
up
front-up
front-level
front-down
×10-4
1.5
1
0.5
0
4
8
12
[kHz]

WB model

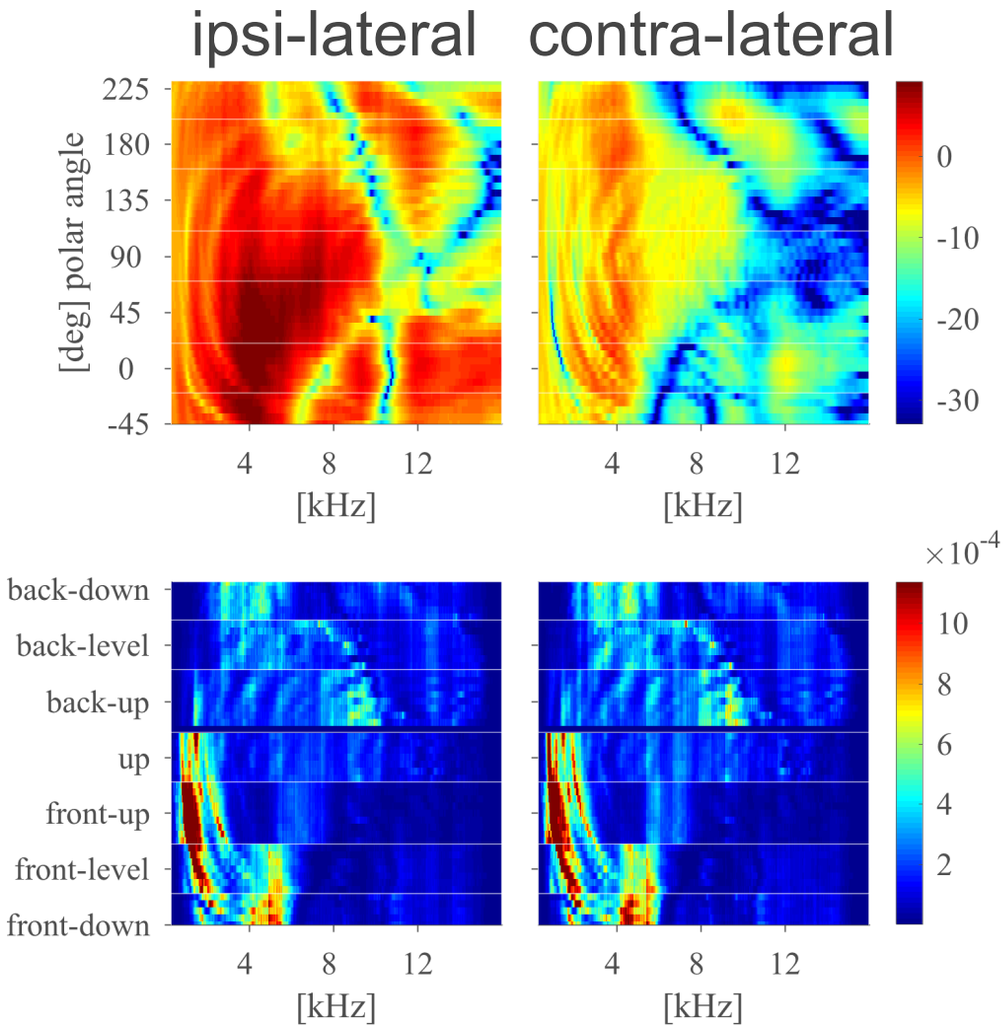
ipsi-lateral
contra-lateral
[deg] polar angle
225
180
135
90
45
0
-45
4
8
12
[kHz]
0
-10
-20
-30
back-down
back-level
back-up
up
front-up
front-level
front-down
×10-4
10
8
6
4
2

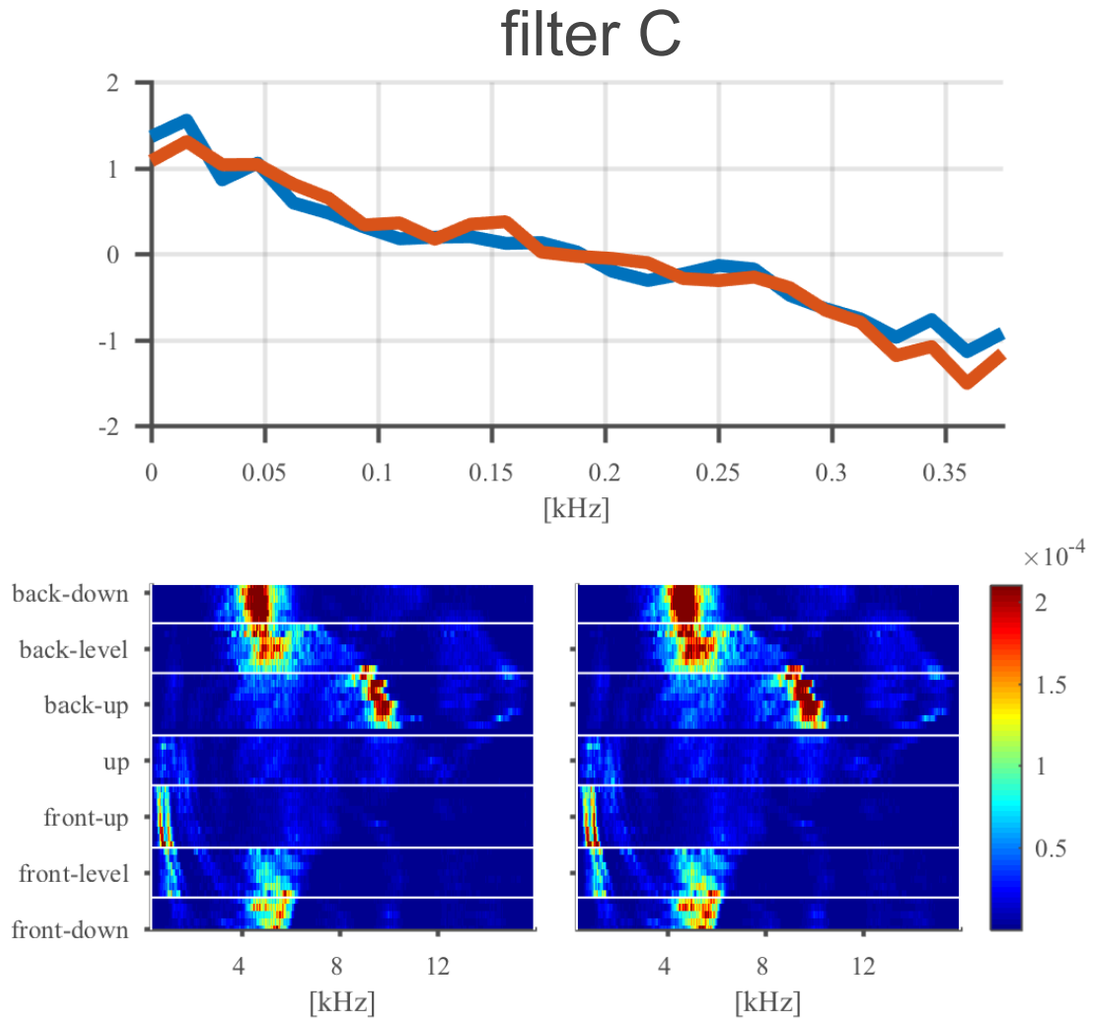
filter C
2
1
0
-1
-2
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
[kHz]
back-down
back-level
back-up
up
front-up
front-level
front-down
×10-4
2
1.5
1
0.5
4
8
12
[kHz]

Conclusion

Summary and conclusion

- (Spatial) audio feature discovery using neural network explanation techniques
- Rudimentary models seem to learn cues similar to ones reported in listening test experiments
- Deep Taylor Decomposition seems useful for discovering/visualizing audio features
- Results to be confirmed using a perceptually-grounded front-end
- Lots of future work…!

Thank you!