

Modeling the Importance of Faces In Natural Images

Bin Jin¹, Gökhan Yildirim¹, Cheryl Lau¹, Appu Shaji¹
Maria V. Ortiz Segovia², Sabine Süsstrunk¹

1. École Polytechnique Fédérale de Lausanne

2. Océ




ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE



A CANON COMPANY

Motivation




Equal importance?




Equal importance?




Equal importance?



Equal importance?



Equal importance?



Equal importance?

The importance of faces **varies** in different images!



Equal importance?

The importance of faces **varies** in different images!

1. Rudoy, D., Goldman, D. B., Shechtman, E., and Zelnik-Manor, L., “Learning video saliency from human gaze using candidate selection,” CVPR, 2013.
2. Marat, S., Guyader, N., and Pellerin, D., “Gaze prediction improvement by adding a face feature to a saliency model,” Recent advances in signal processing , 2009.
3. Judd, T., Ehinger, K., Durand, F., and Torralba, A., “Learning to predict where humans look,” ICCV, 2106–2113 (2009).



Outline

Outline

- Crowd-sourcing experiment

Outline

- Crowd-sourcing experiment
- Face-importance model

Outline

- Crowd-sourcing experiment
- Face-importance model
- Salient-object detection

Outline

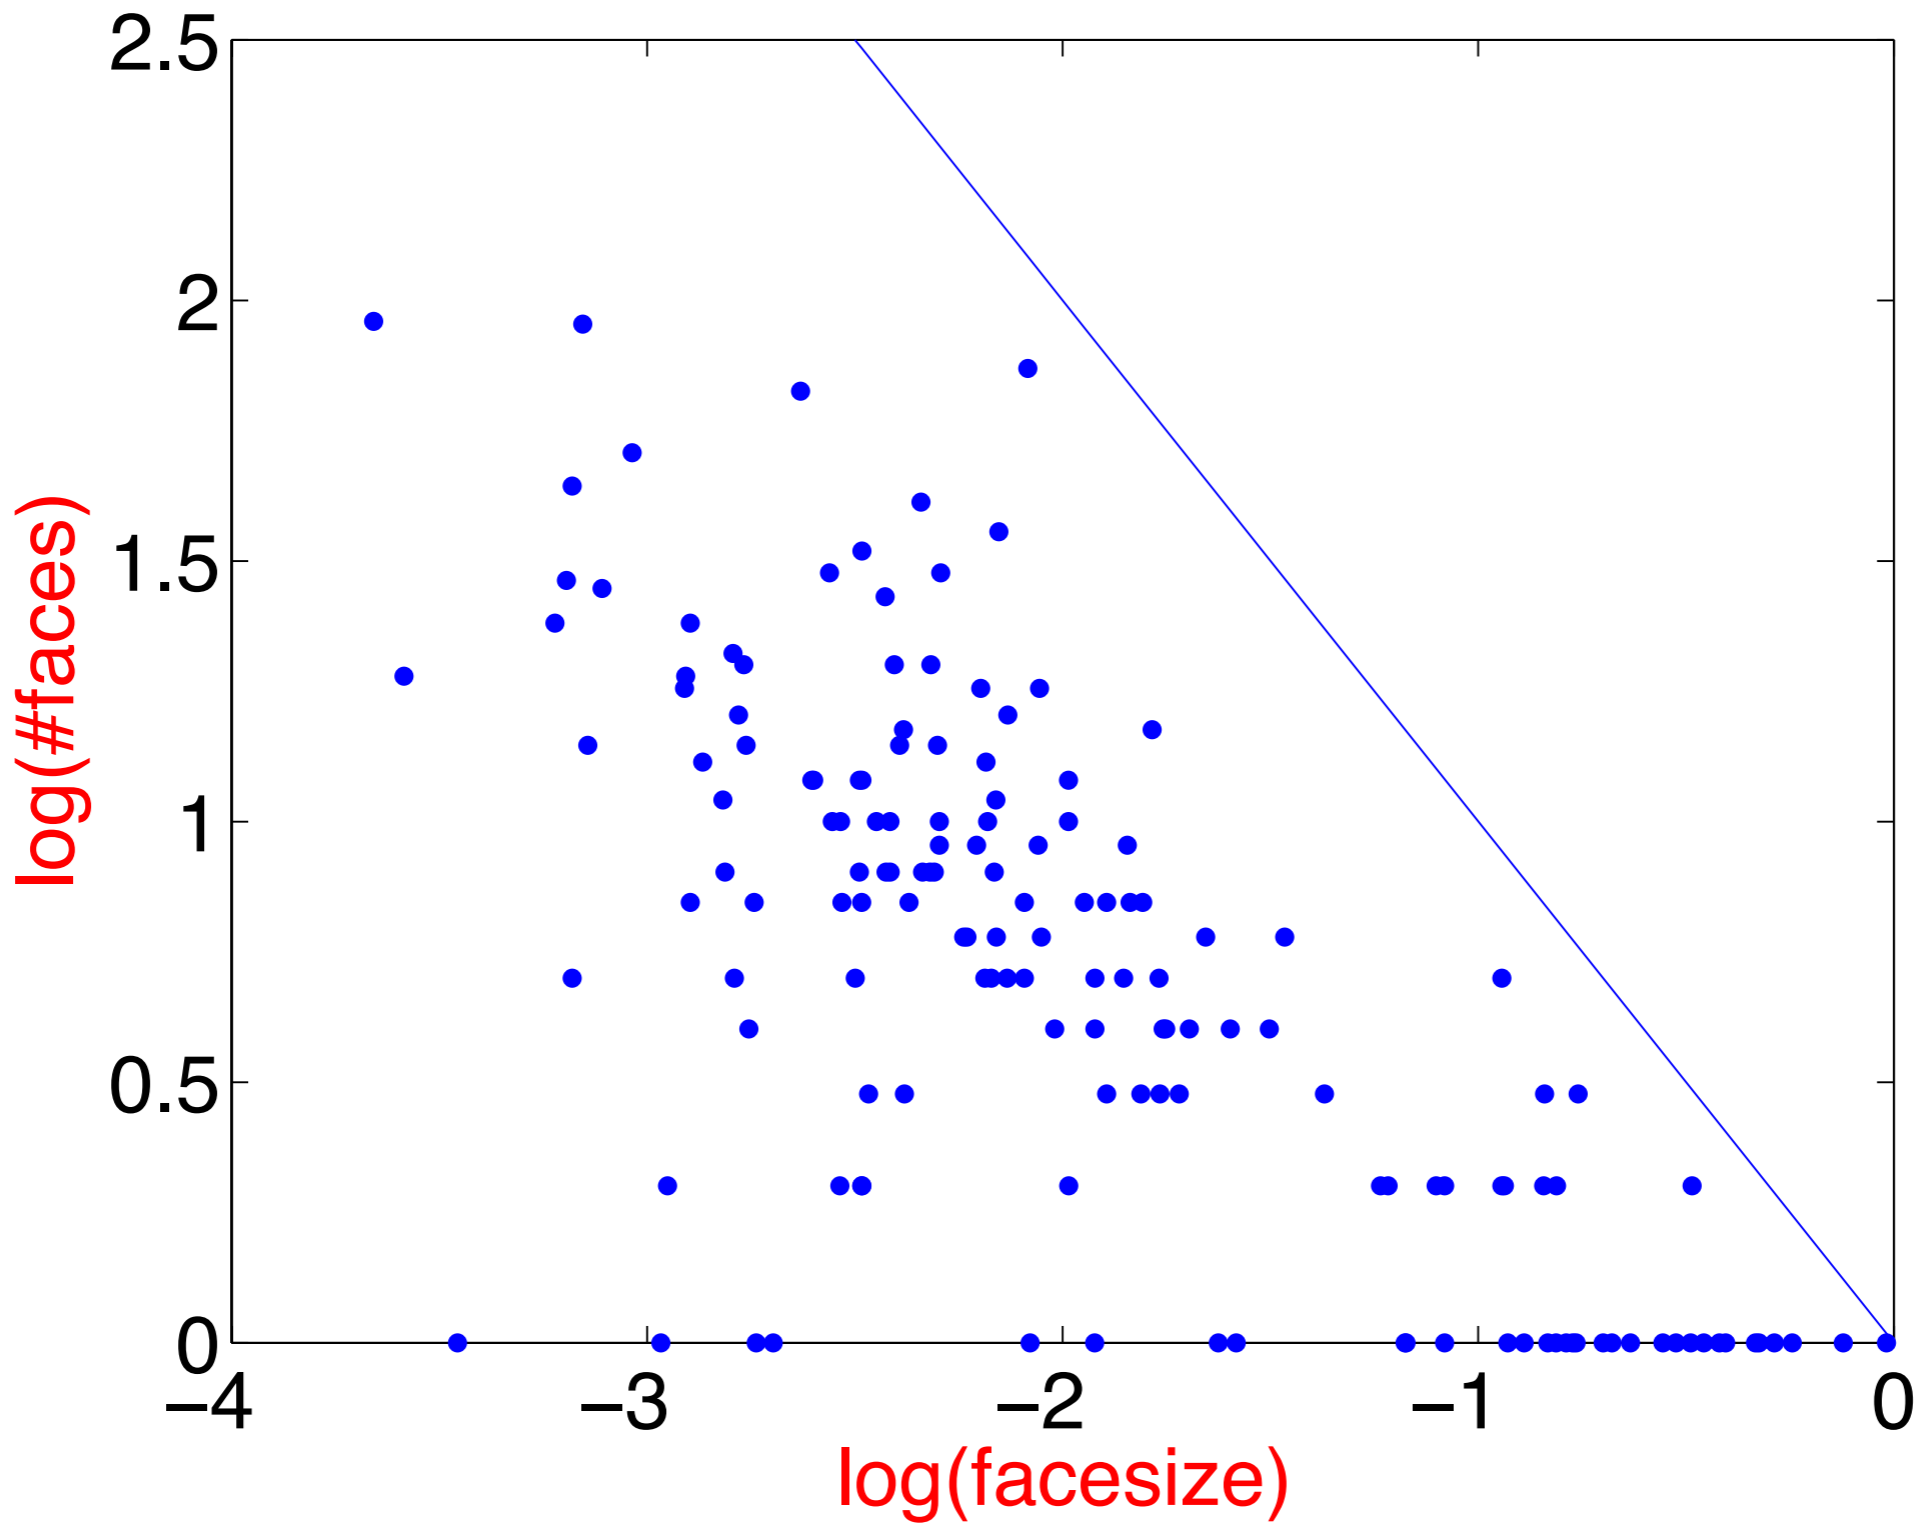
- Crowd-sourcing experiment
- Face-importance model
- Salient-object detection
- Conclusion

Outline

- Crowd-sourcing experiment
- Face-importance model
- Salient-object detection
- Conclusion

Dataset

152 face images



Crowd-sourcing experiment

Reward: 0.30

Creation Date: 2014-01-09 09:54:32

Tasks Available: 30

Expiration Date: 2014-01-19 09:54:32

BACK

Description: 50 images for T1

Instructions: For each image, label out the important parts in the image. It could be a face, or eyes or human bodies or even other objects. Draw rectangles on the area that you think is important. Hold left click and move your mouse to draw rectangles.

Label out the important parts of the image

1. Go to website http://ivrgwww.epfl.ch/~bjin/forsaliency/Task1/labelimages_intro.html;
2. Read the discription. Then click on CreateUser button and enter a user name.Start labelling images.
3. You will be shown 50 images with each at a time. Label out the important objects in the image by drawing rectangles.

2. Write your username in below and I will check your results and then pay you:

BACK

Crowd-sourcing experiment

Reward: 0.30

Creation Date: 2014-01-09 09:54:32

Tasks Available: 30

Expiration Date: 2014-01-19 09:54:32

BACK

Description: 50 images for T1

Instructions: For each image, label out the important parts in the image. It could be a face, or eyes or human bodies or even other objects. Draw rectangles on the area that you think is important. Hold left click and move your mouse to draw rectangles.

Label out the important parts of the image

1. Go to website http://ivrgwww.epfl.ch/~bjin/forsaliency/Task1/labelimages_intro.html;
2. Read the discription. Then click on CreateUser button and enter a user name.Start labelling images.
3. You will be shown 50 images with each at a time. Label out the important objects in the image by drawing rectangles.

2. Write your username in below and I will check your results and then pay you:

BACK

~ 30 observers per image

Crowd-sourcing experiment



Label out the important objects in the image

Label out the important objects in the image by drawing rectangles. Draw a rectangle on each important objects. Hold left click and move your mouse to draw rectangles. Press key z to delete the last rectangle you draw. Click on next image to go to next image.







next image

Crowd-sourcing results




Crowd-sourcing results



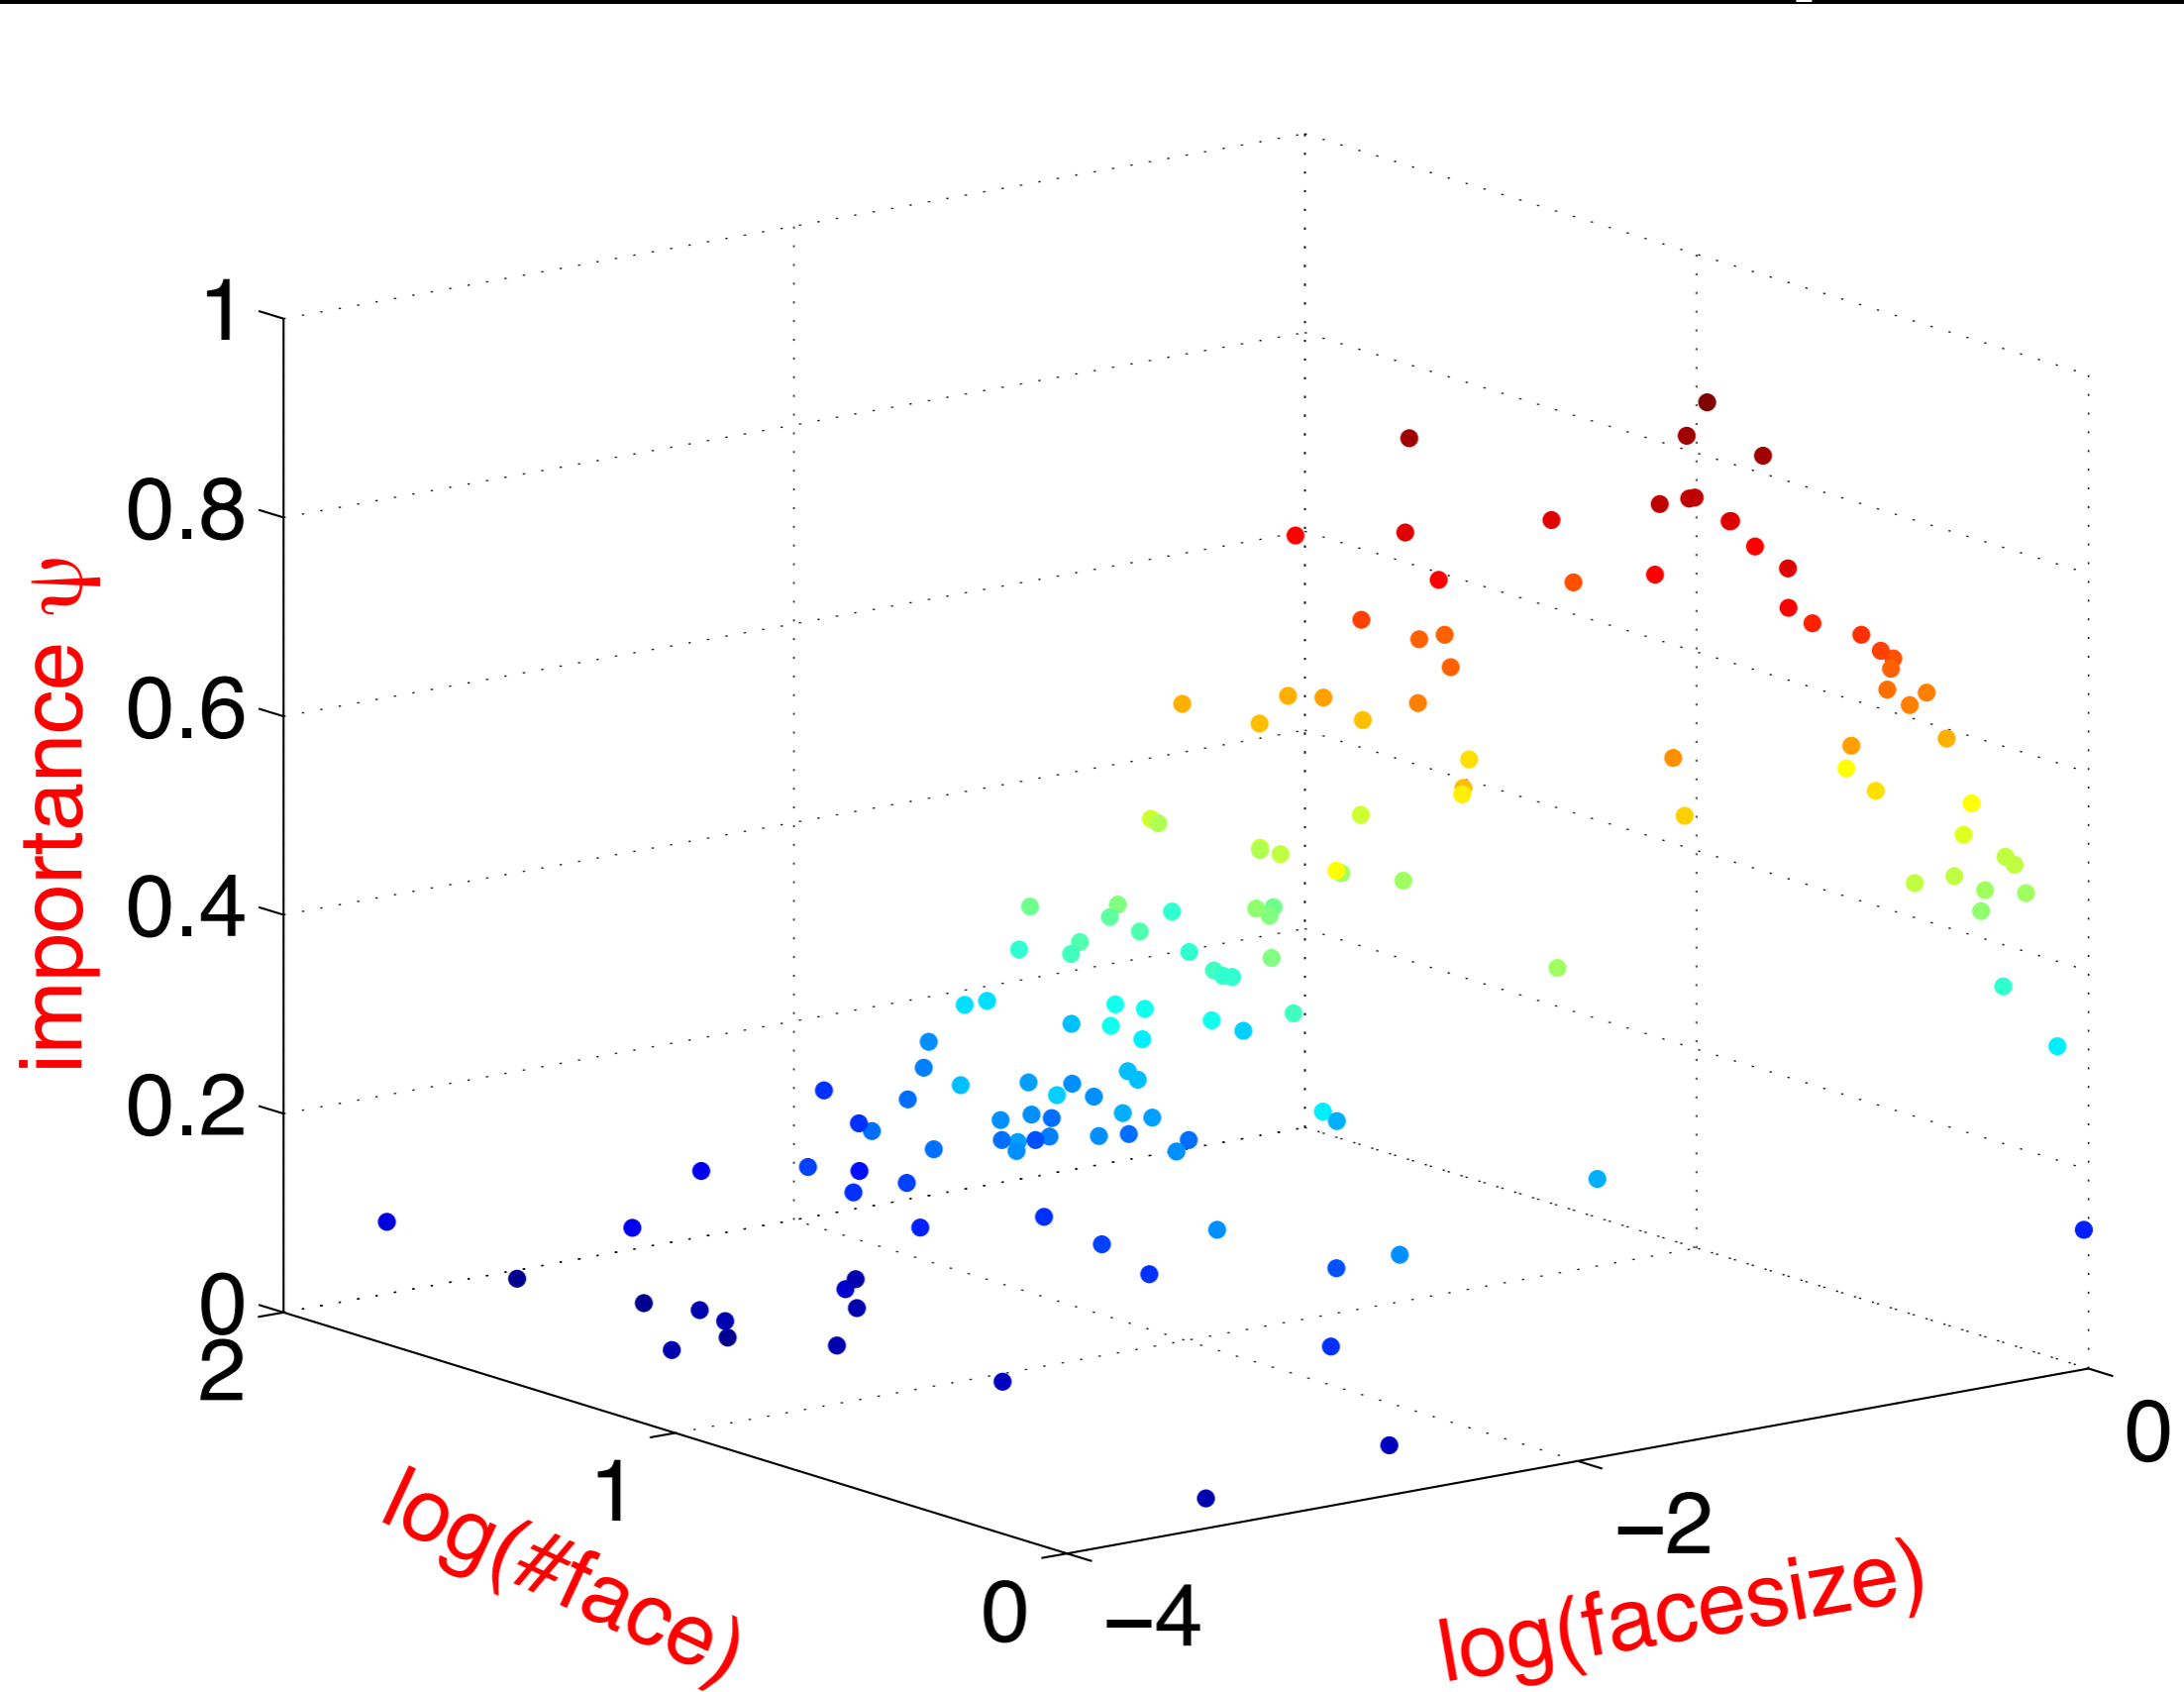
Outline

- Crowd-sourcing experiment
- **Face-importance model**
- Salient-object detection
- Conclusion


ID distribution of face importance



Joint distribution of face importance

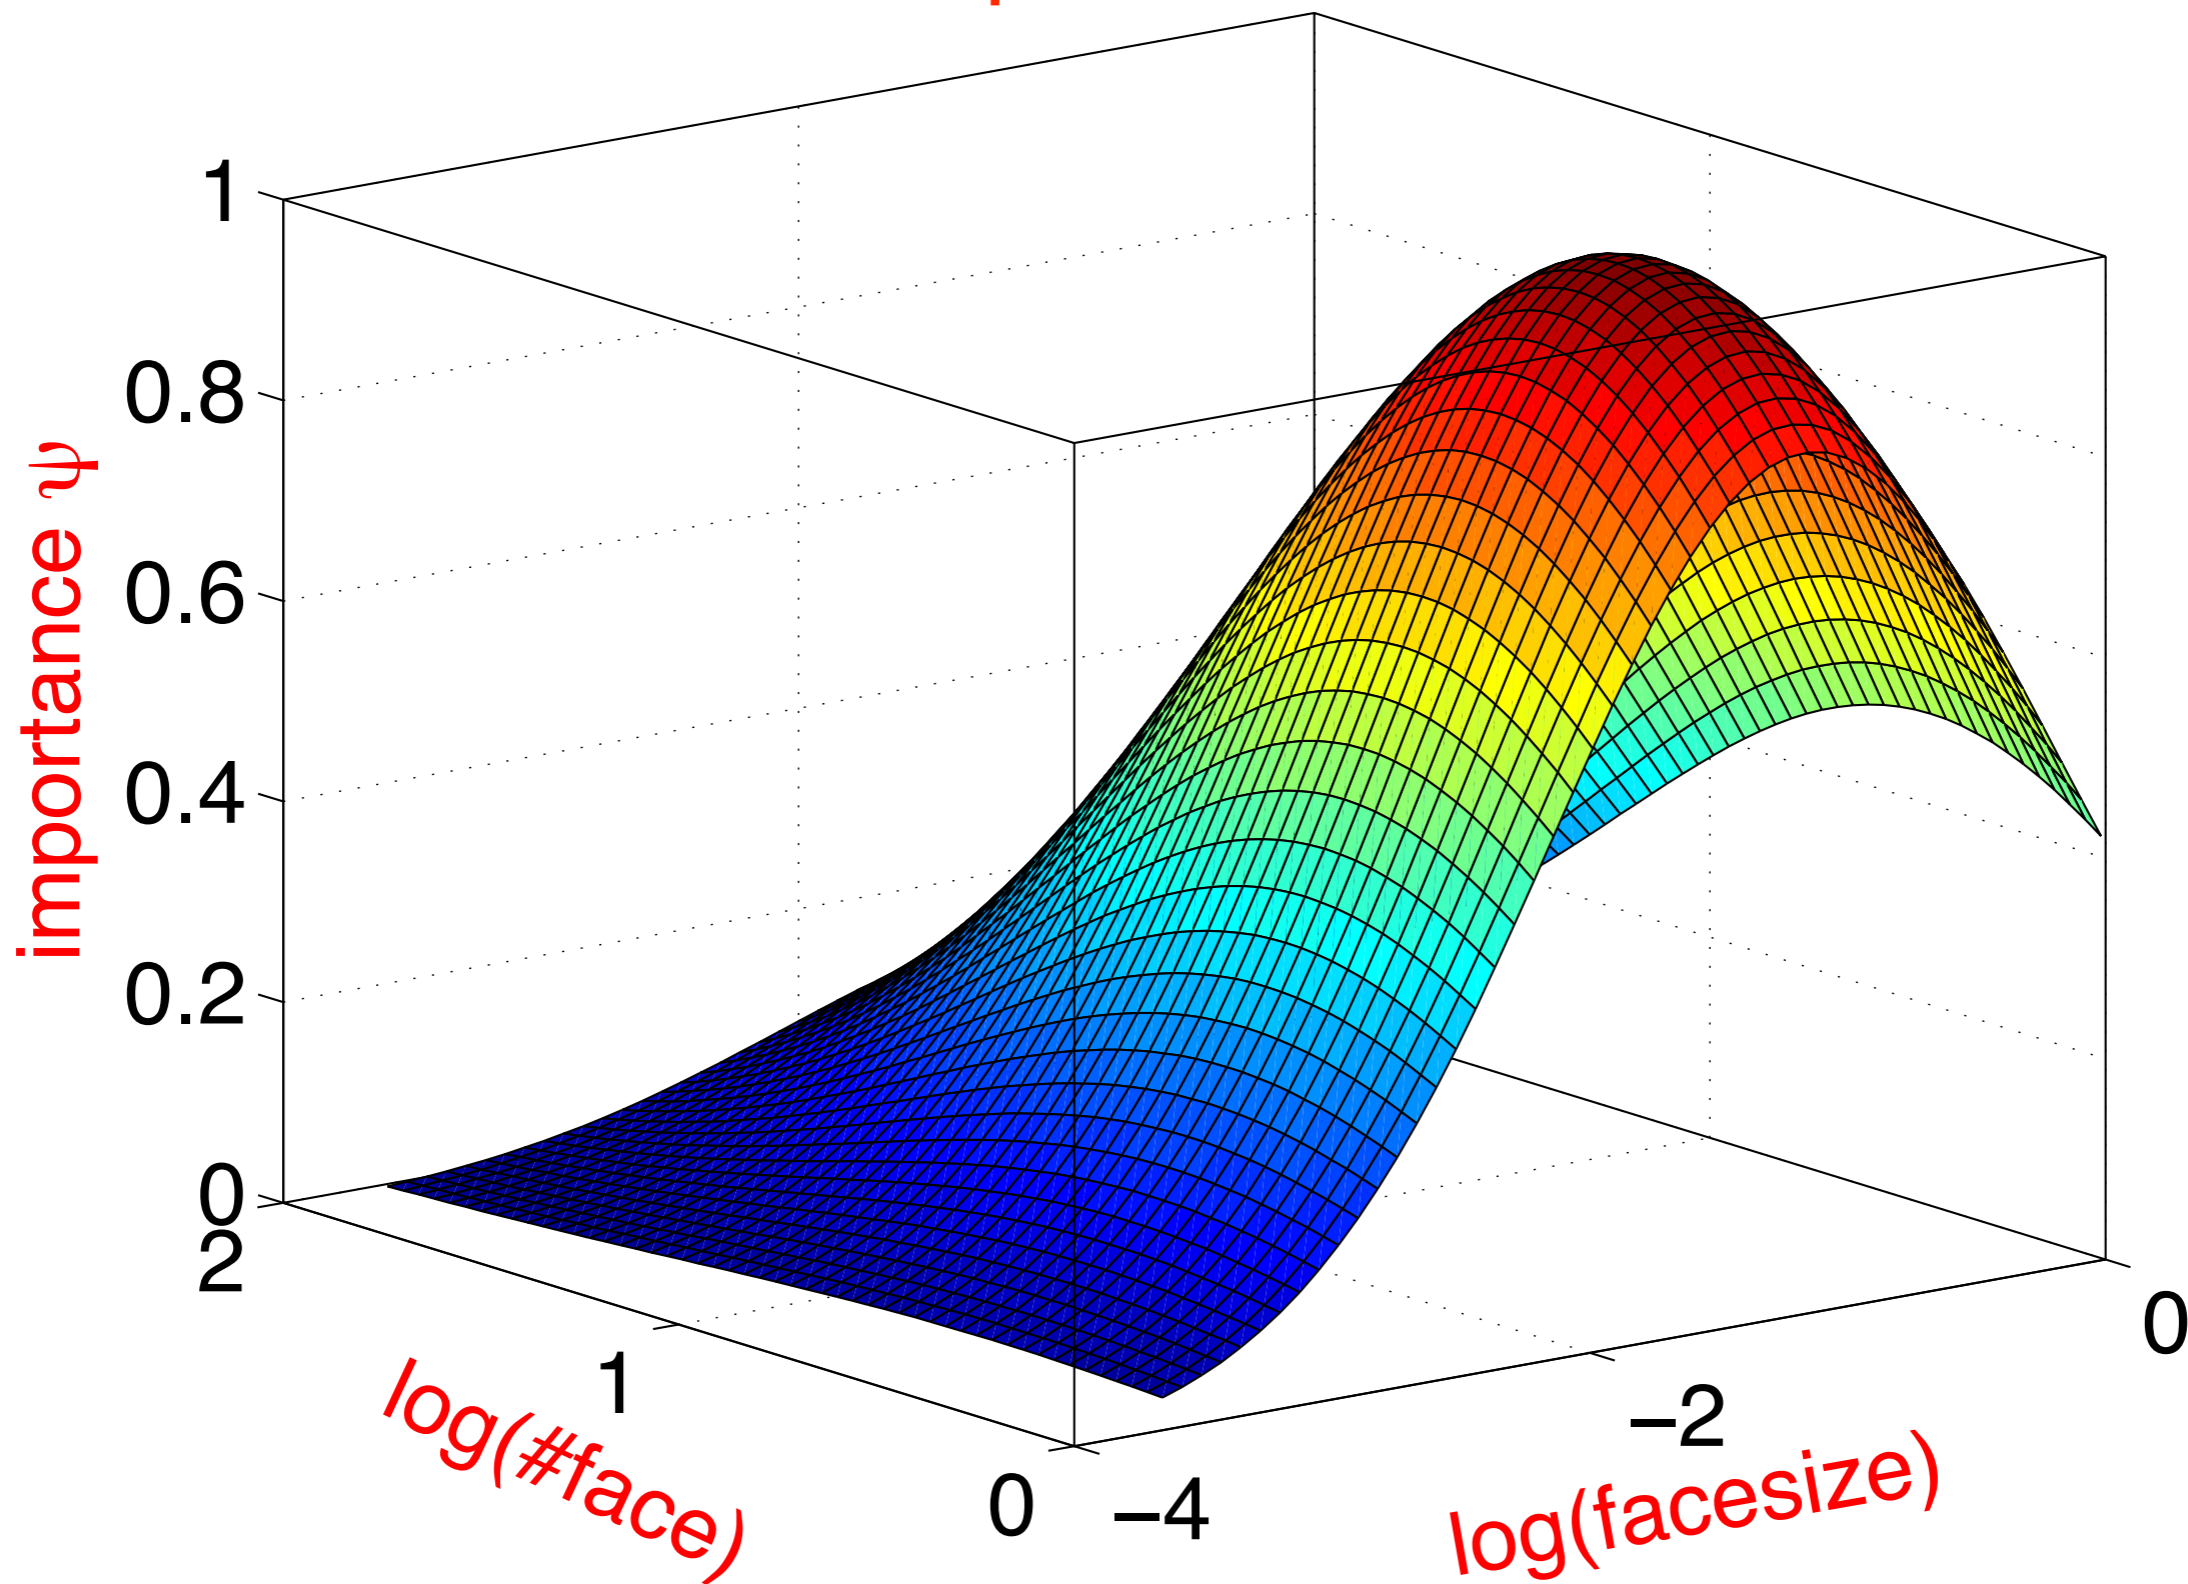


Face importance model




Face importance model

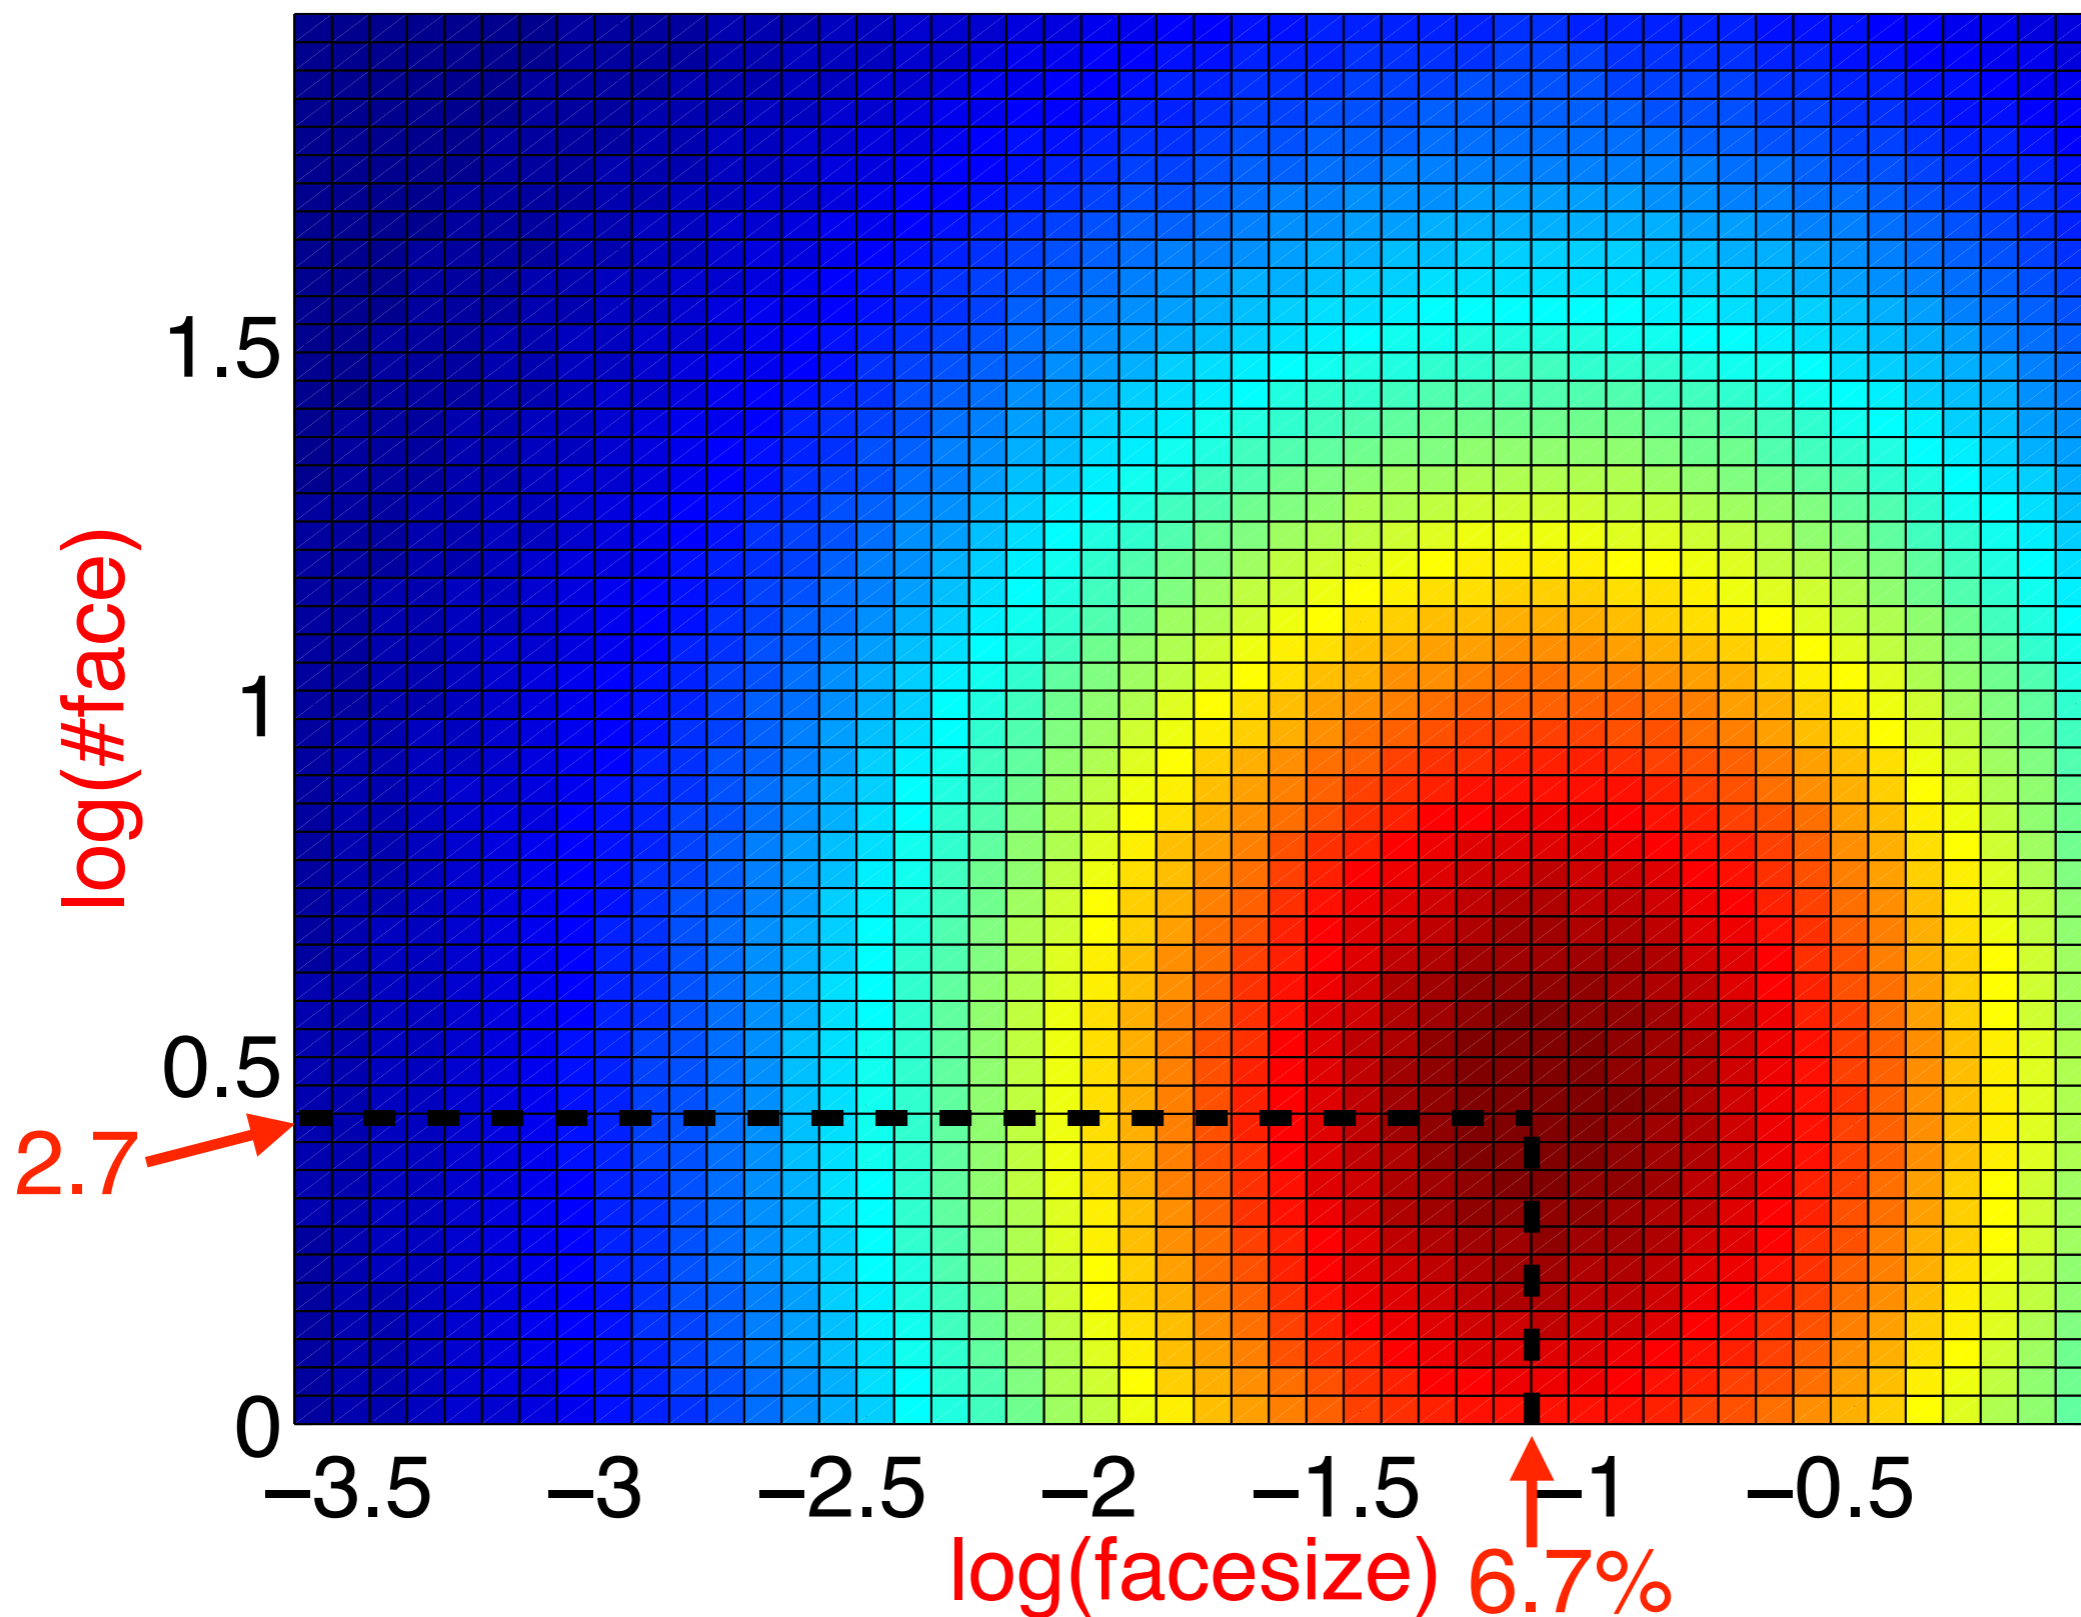
R-square: 0.8567



Heatmap



Heatmap



Outline

- Crowd-sourcing experiment
- Face-importance model
- **Salient-object detection**
- Conclusion

Salient-object detection

original image



saliency map



Saliency algorithm

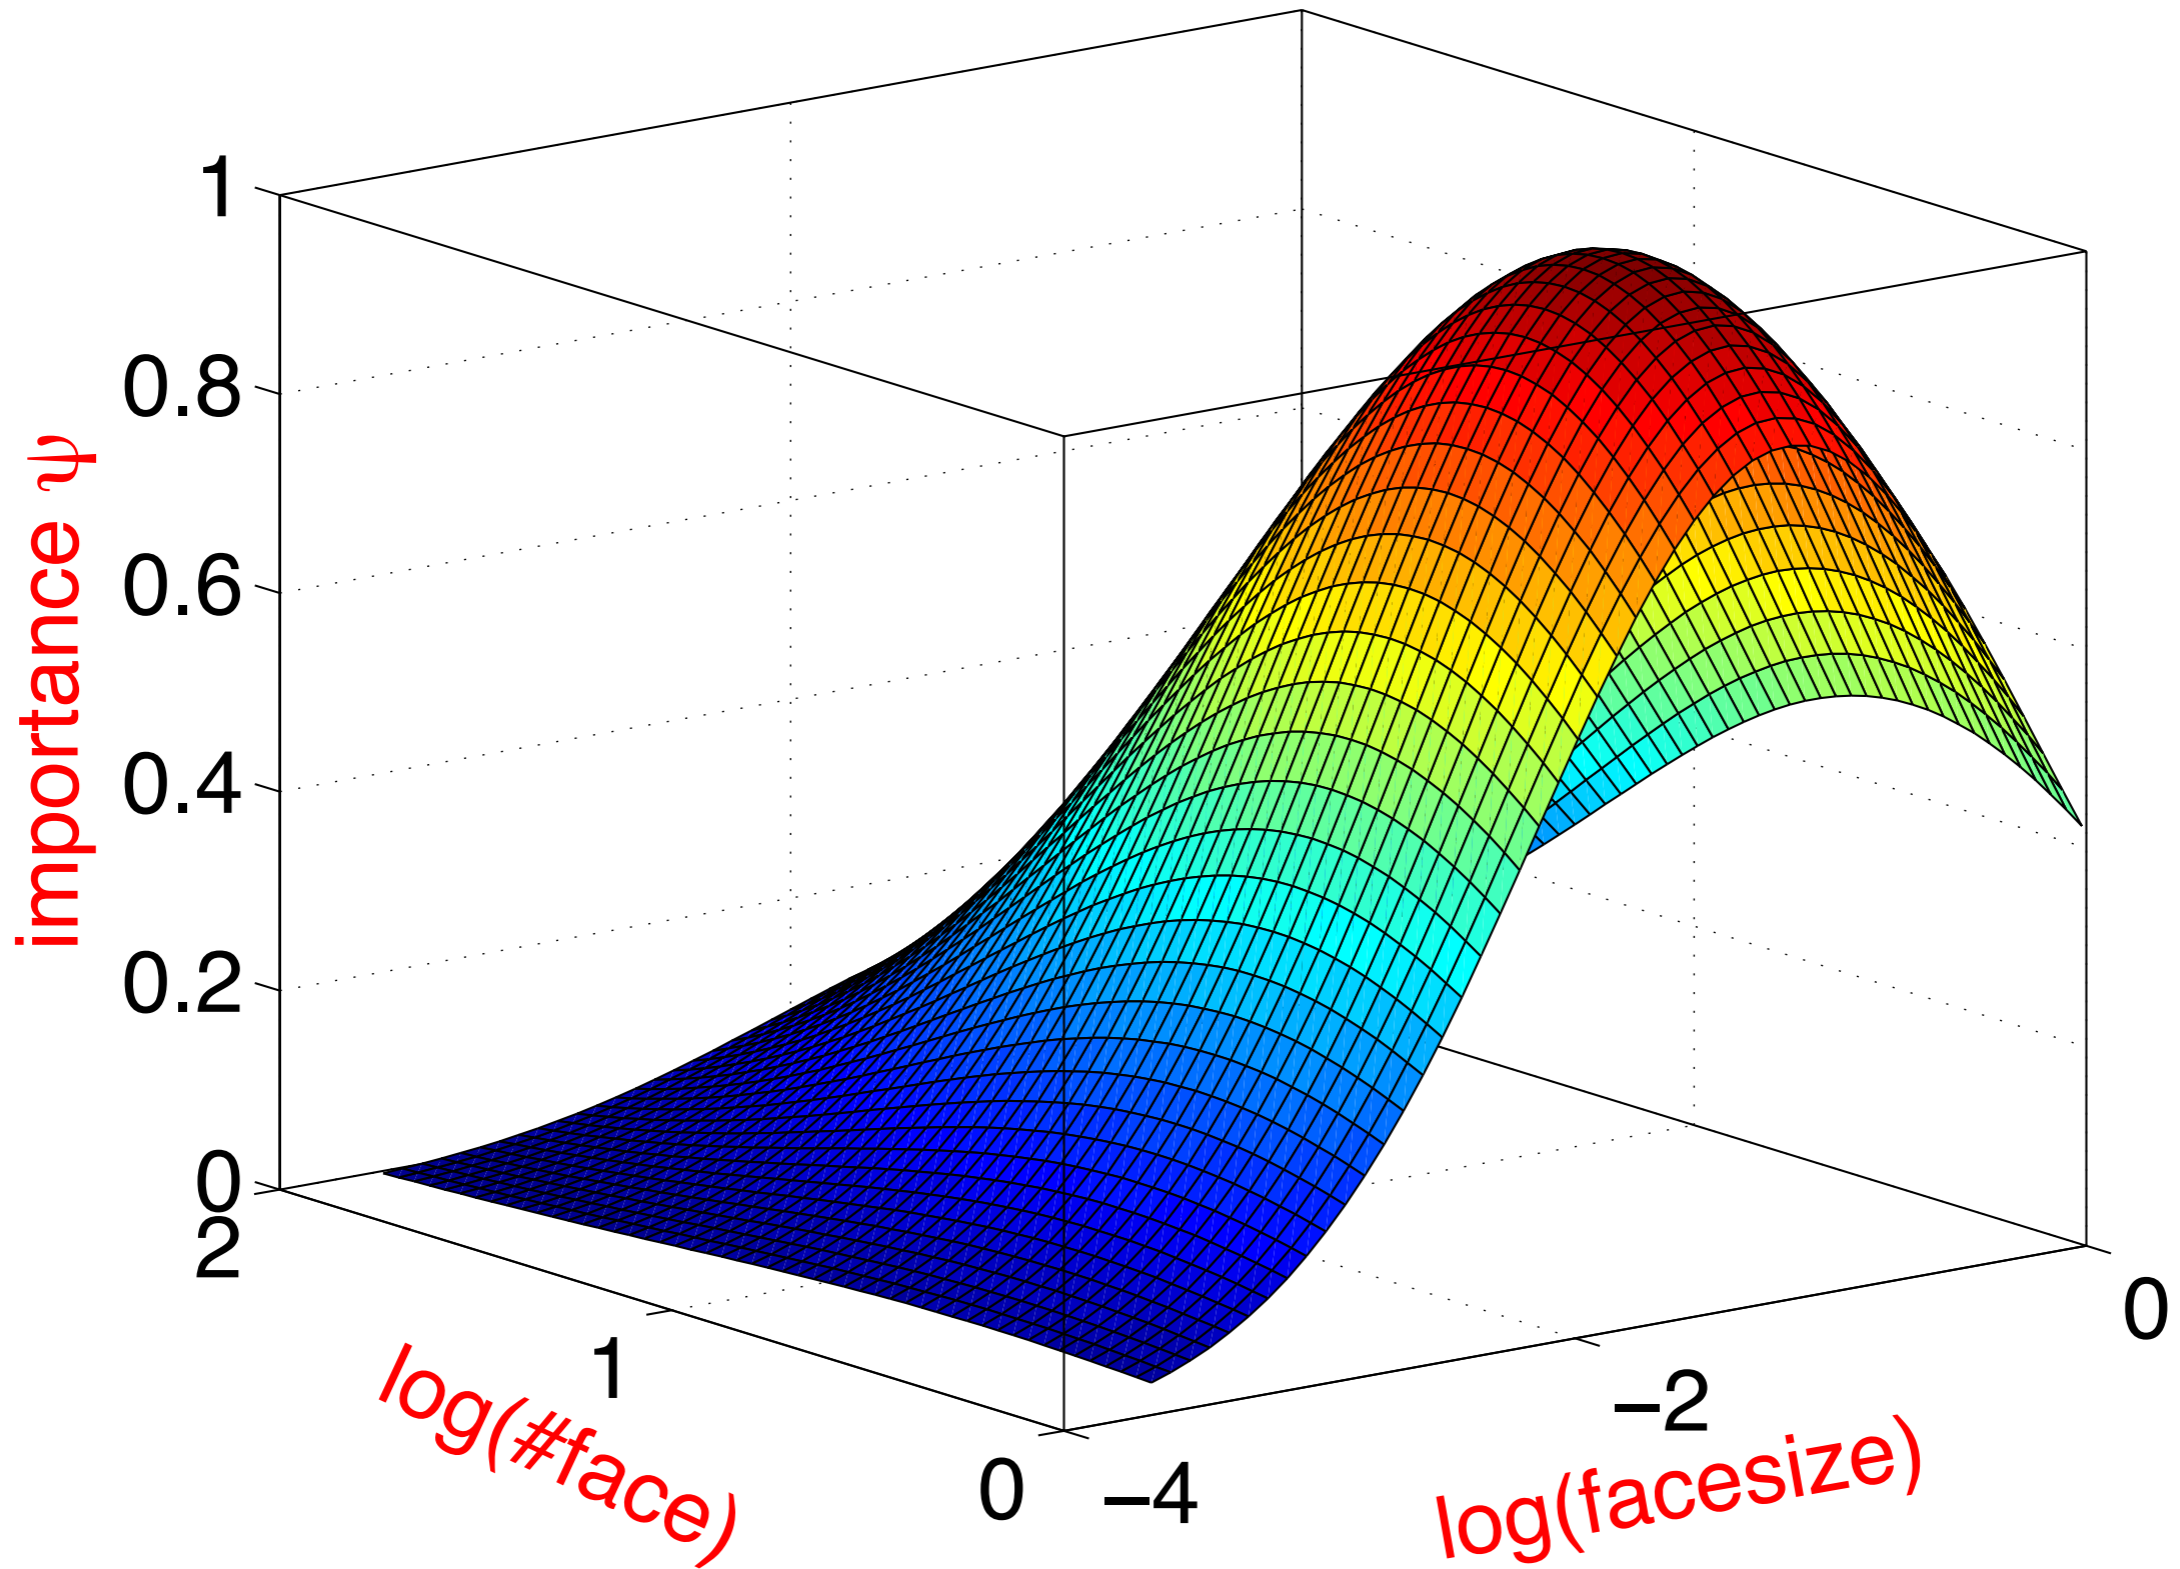


Saliency algorithm





input

Saliency algorithm




Saliency algorithm



Saliency algorithm



Saliency algorithm



face-importance map

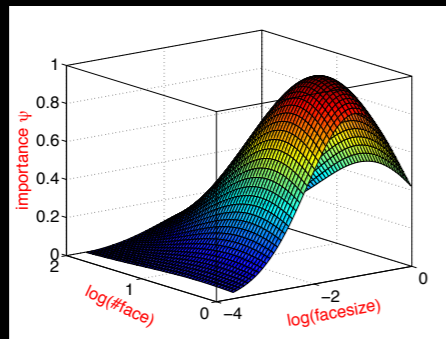


input

Saliency algorithm



Saliency algorithm



face-importance map




contrast-based saliency map

input



Saliency algorithm



face-importance map

contrast-based saliency map


input



Saliency algorithm



Saliency algorithm



face-importance map



input

contrast-based saliency map



final saliency map



Face importance map



Face importance map



Face importance map




Face importance map



Face importance map




Face importance map




Face importance map




Face importance map






Face importance map






Face importance map



Face importance map



Face importance map




Contrast-based saliency map



Contrast-based saliency map



Combination



face-importance map




input

contrast-based saliency map



final saliency map

Combination




face-importance map




input

contrast-based saliency map



final saliency map

Combination




face-importance map



input

contrast-based saliency map




final saliency map


Evaluation dataset




Evaluation dataset



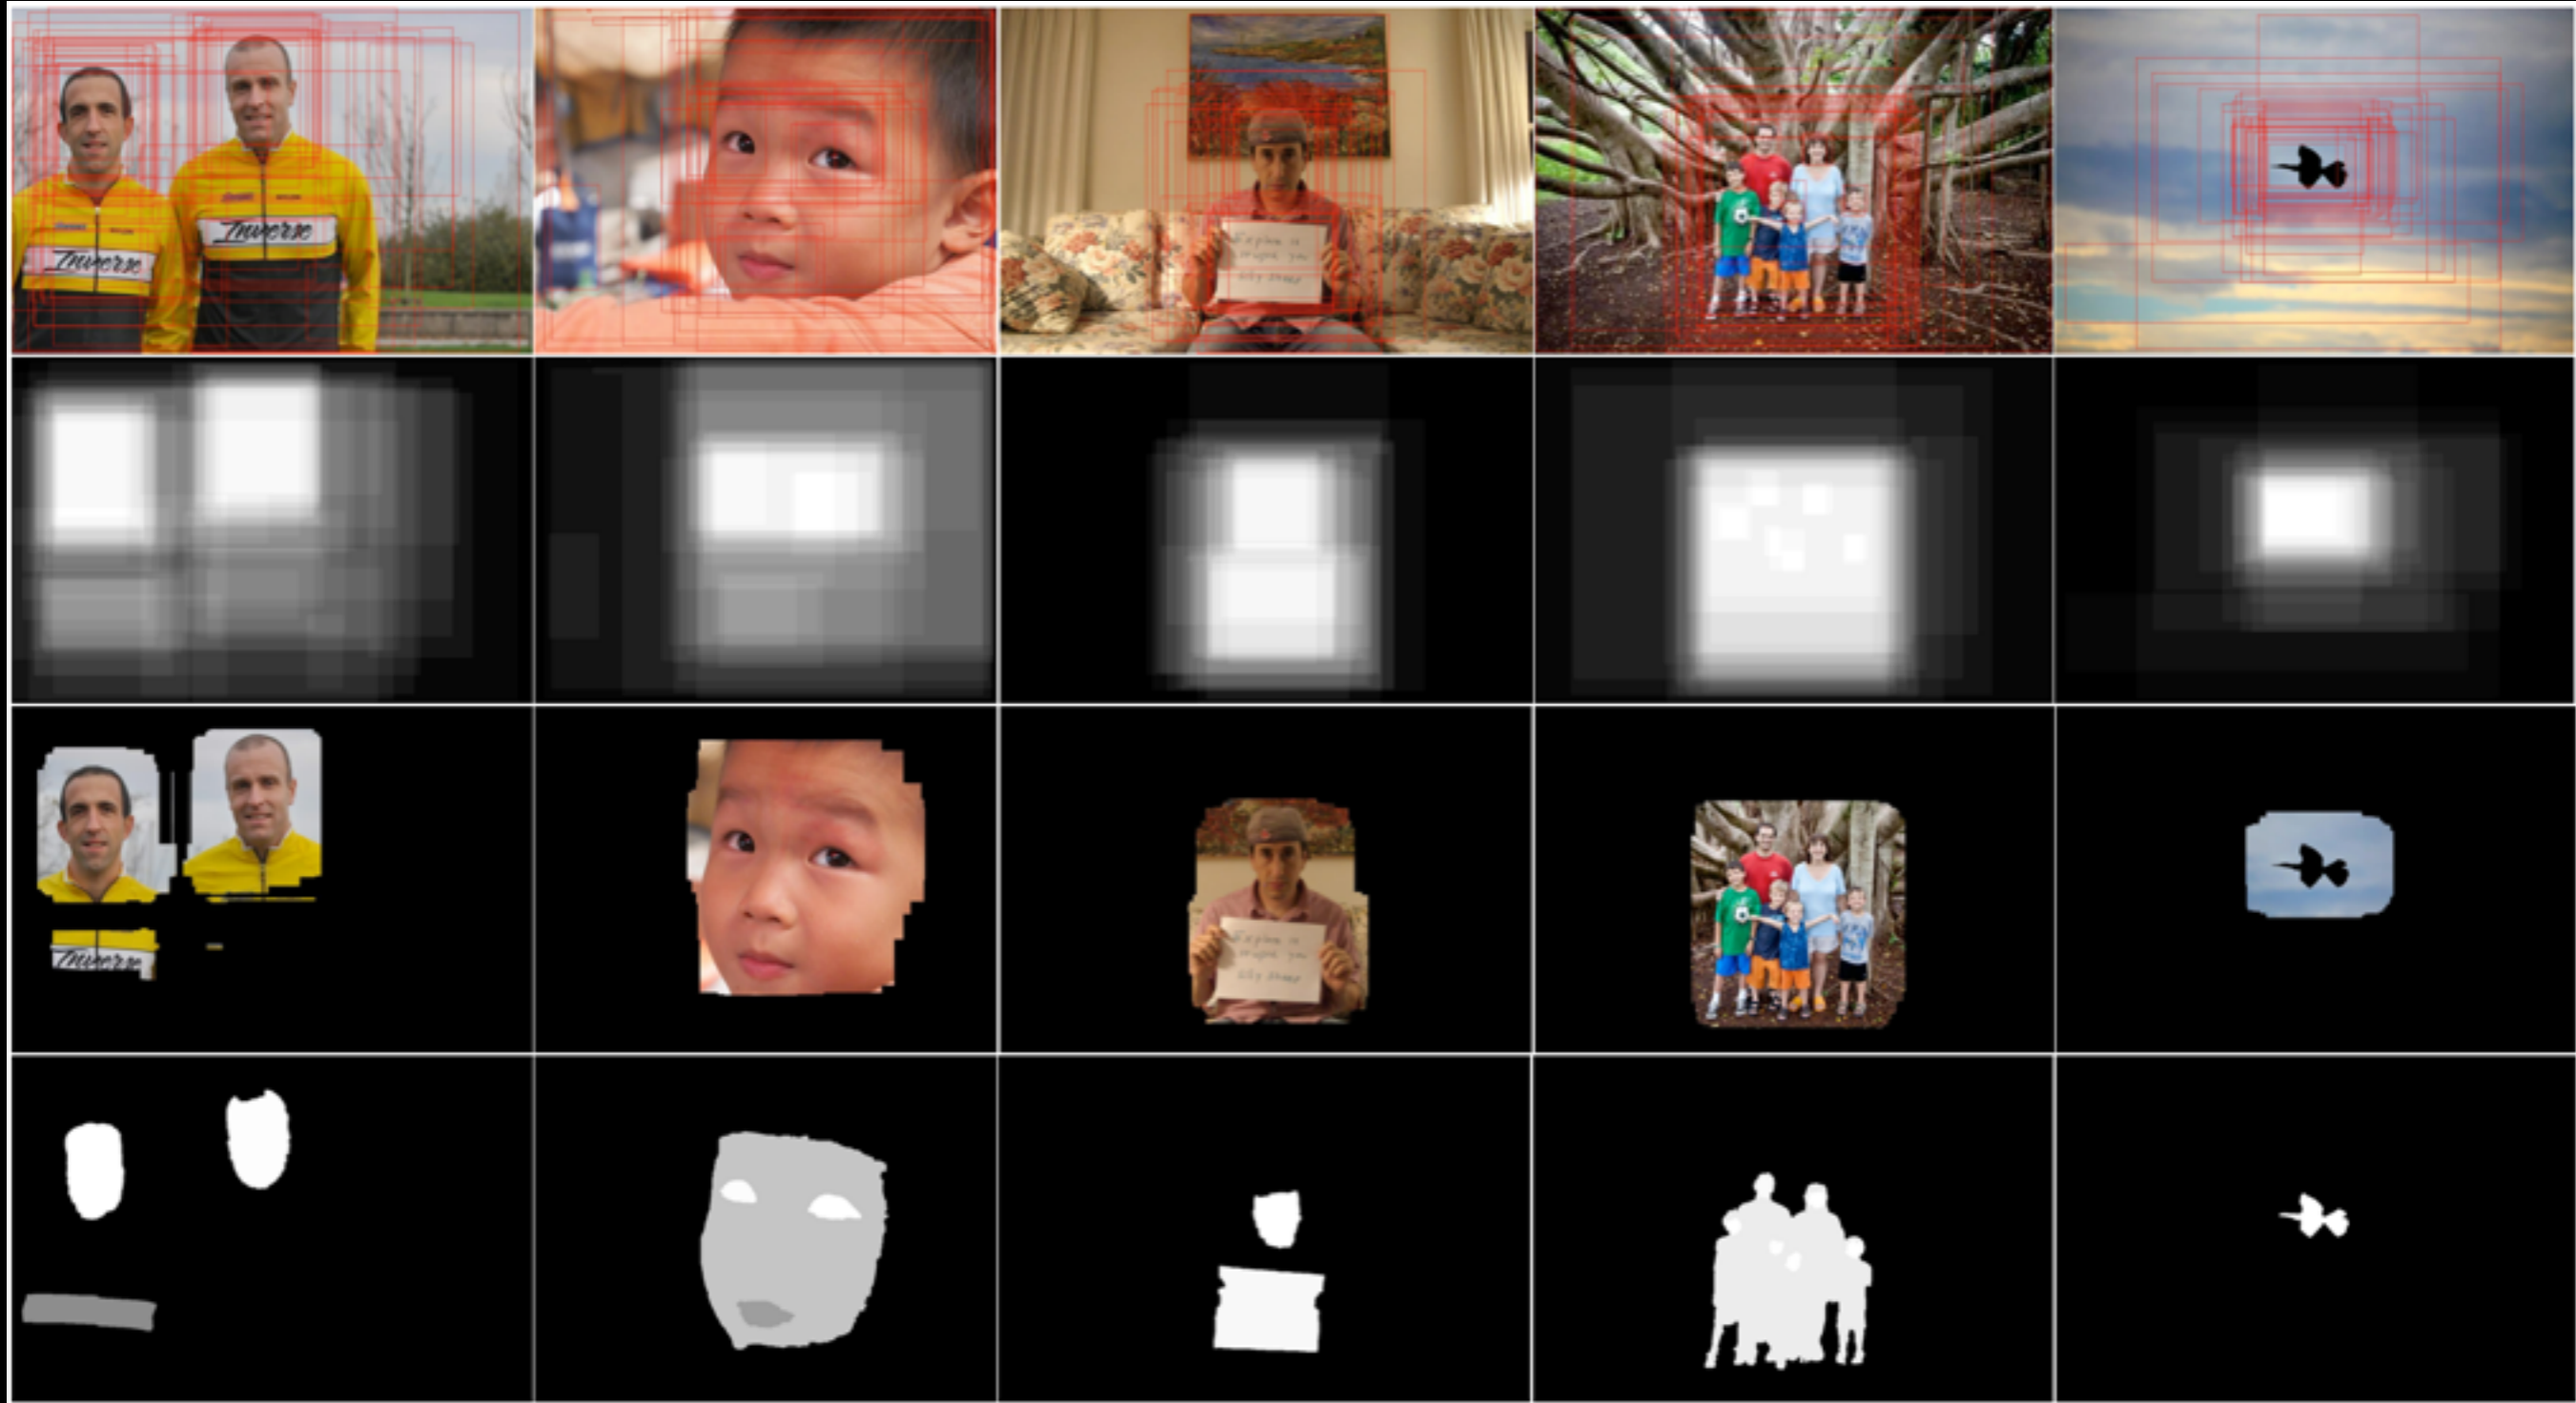
Ground-truth




Ground-truth




Ground-truth



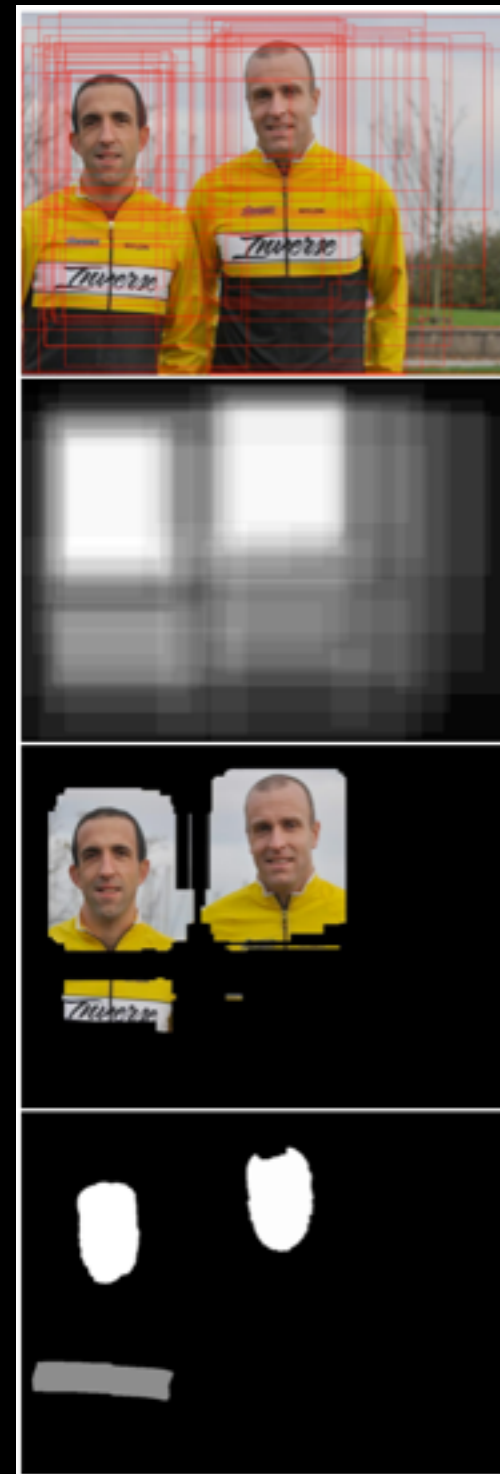
Ground-truth




Ground-truth



Ground-truth




Ground-truth



Ground-truth

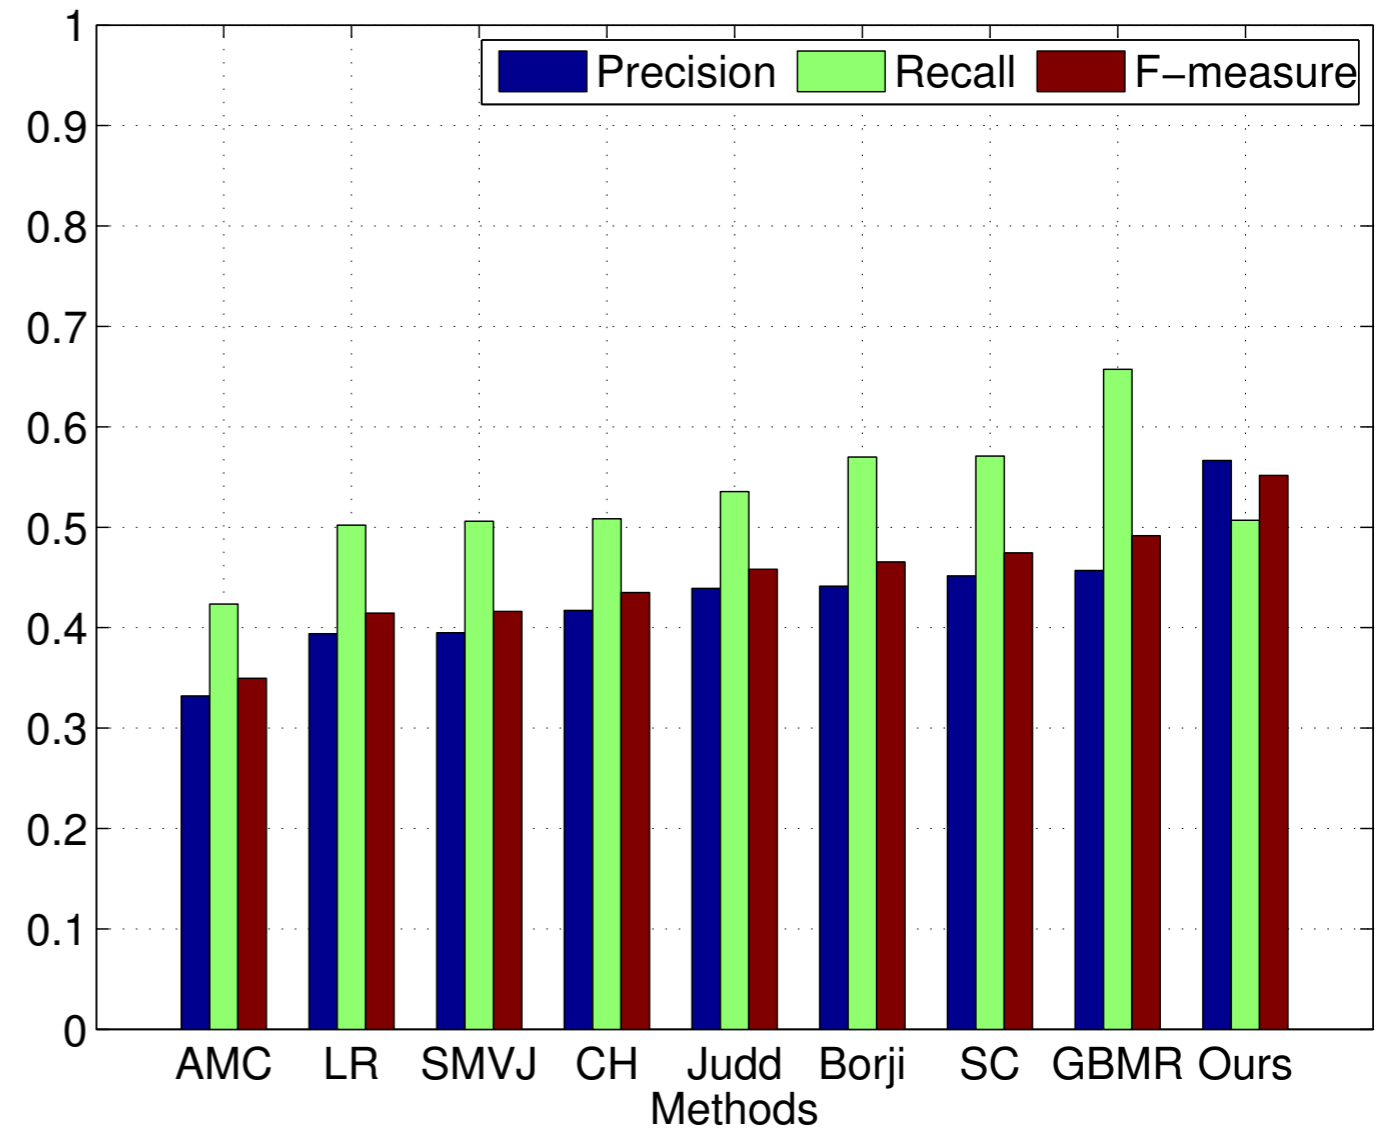
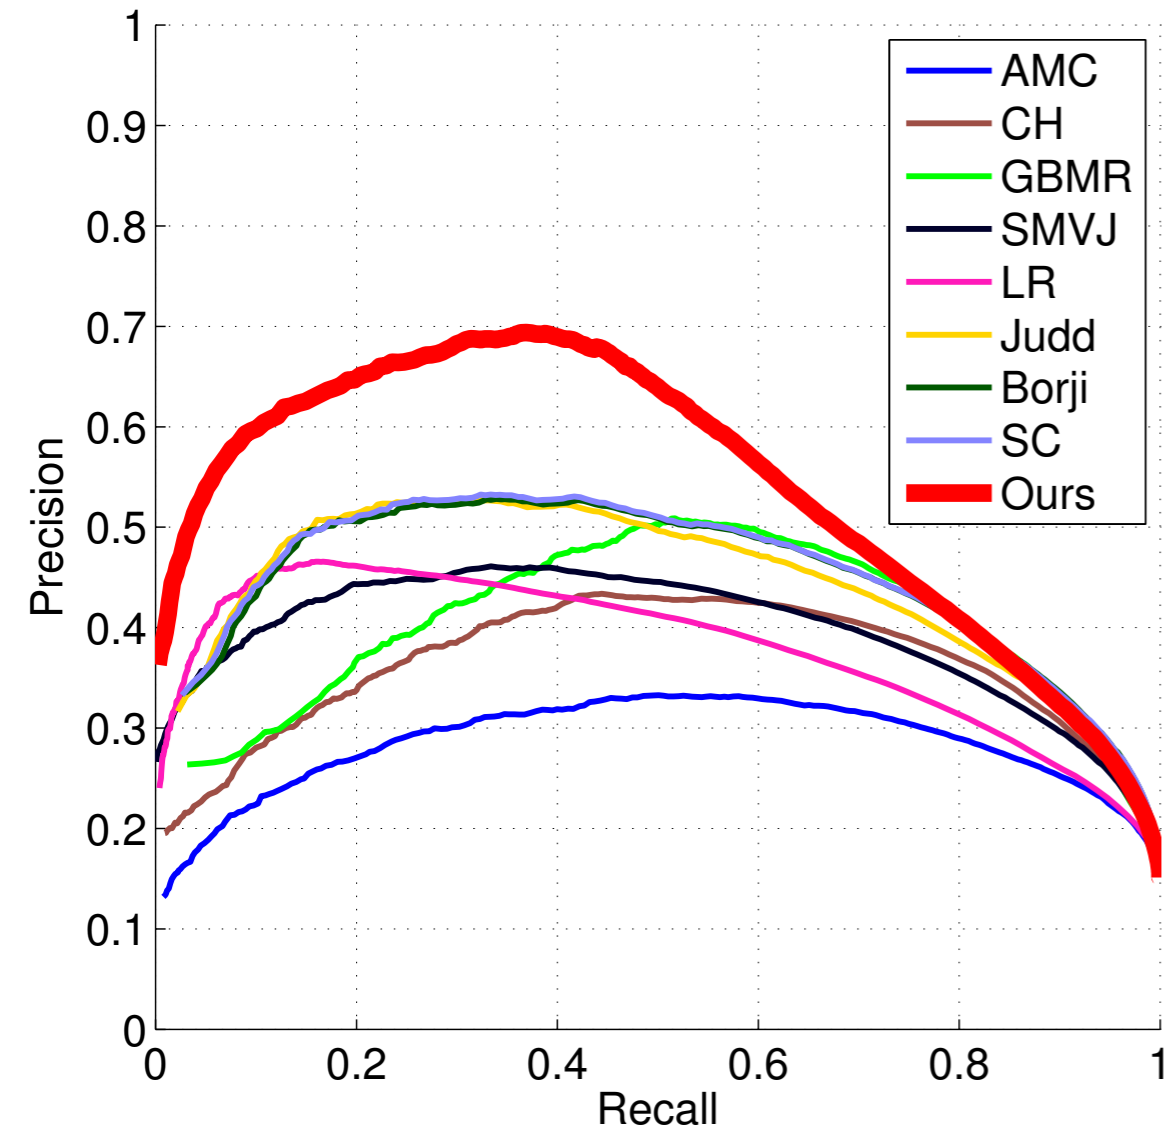


Ground-truth


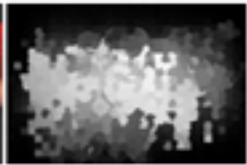
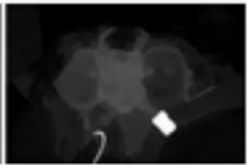



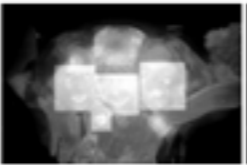
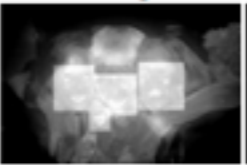
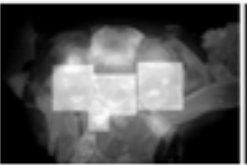



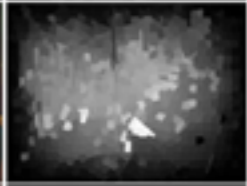
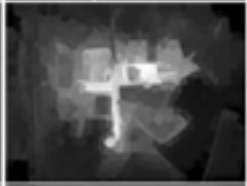












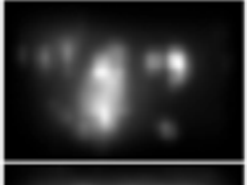







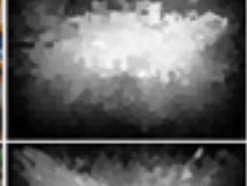
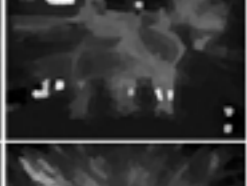



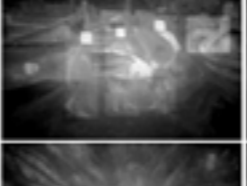
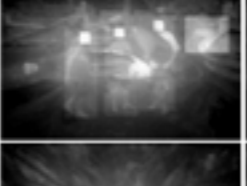







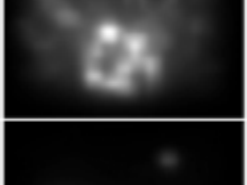

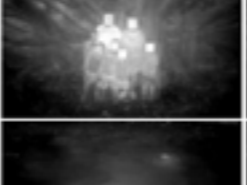
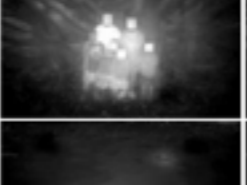
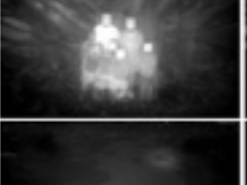



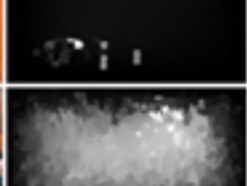
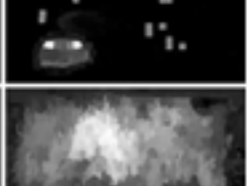

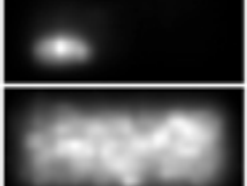
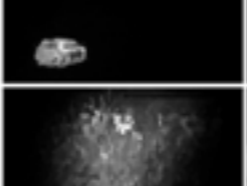
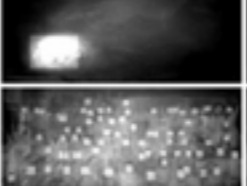
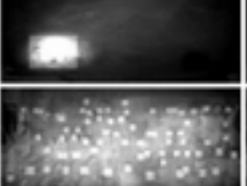




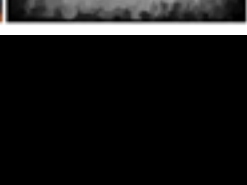
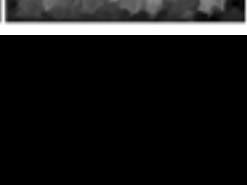










Results


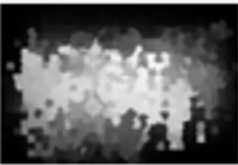
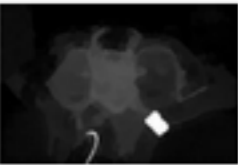
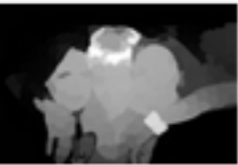


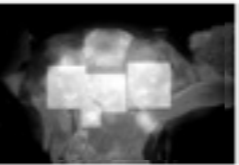
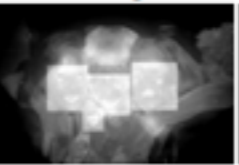
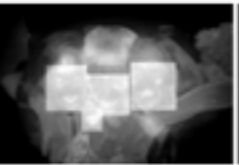









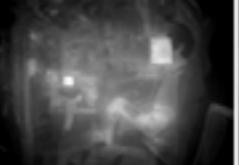







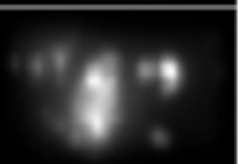








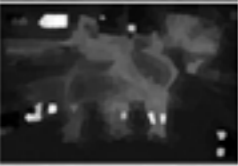






















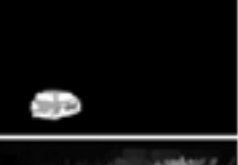
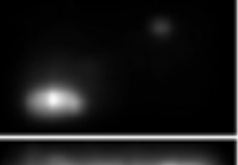





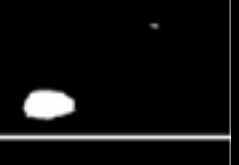



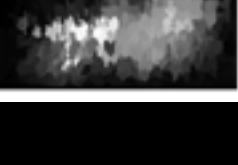







PR curves



Results

Input	AMC	CH	GBMR	SMVJ	LR	Judd	Borji	SC	Ours	Ground truth
										
										
										
										
										
										
										

Results

Input	AMC	CH	GBMR	SMVJ	LR	Judd	Borji	SC	Ours	Ground truth
										
										
										
										
										
										
										

Results



	LR	Judd	Borji	SC	Ours	Ground truth

Results





LR	Judd	Borji	SC	Ours	Ground truth

Results




LR




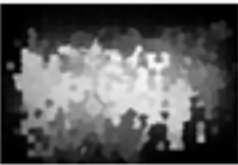
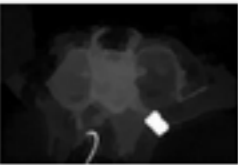
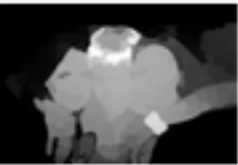


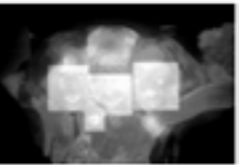
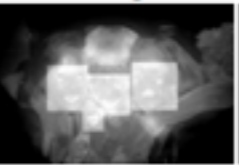
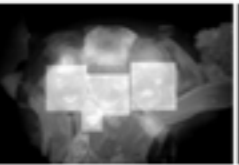

















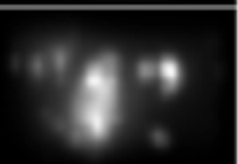








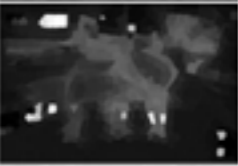






















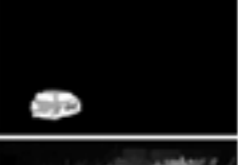
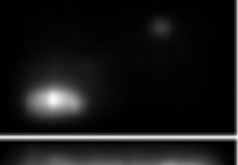





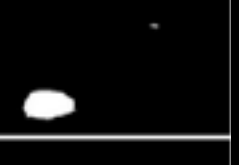



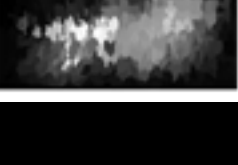







Results



LR



Results

Input	AMC	CH	GBMR	SMVJ	LR	Judd	Borji	SC	Ours	Ground truth
										
										
										
										
										
										
										

Results



	LR	Judd	Borji	SC	Ours	Ground truth


Results



LR	Judd	Borji	SC	Ours	Ground truth




Results




Results



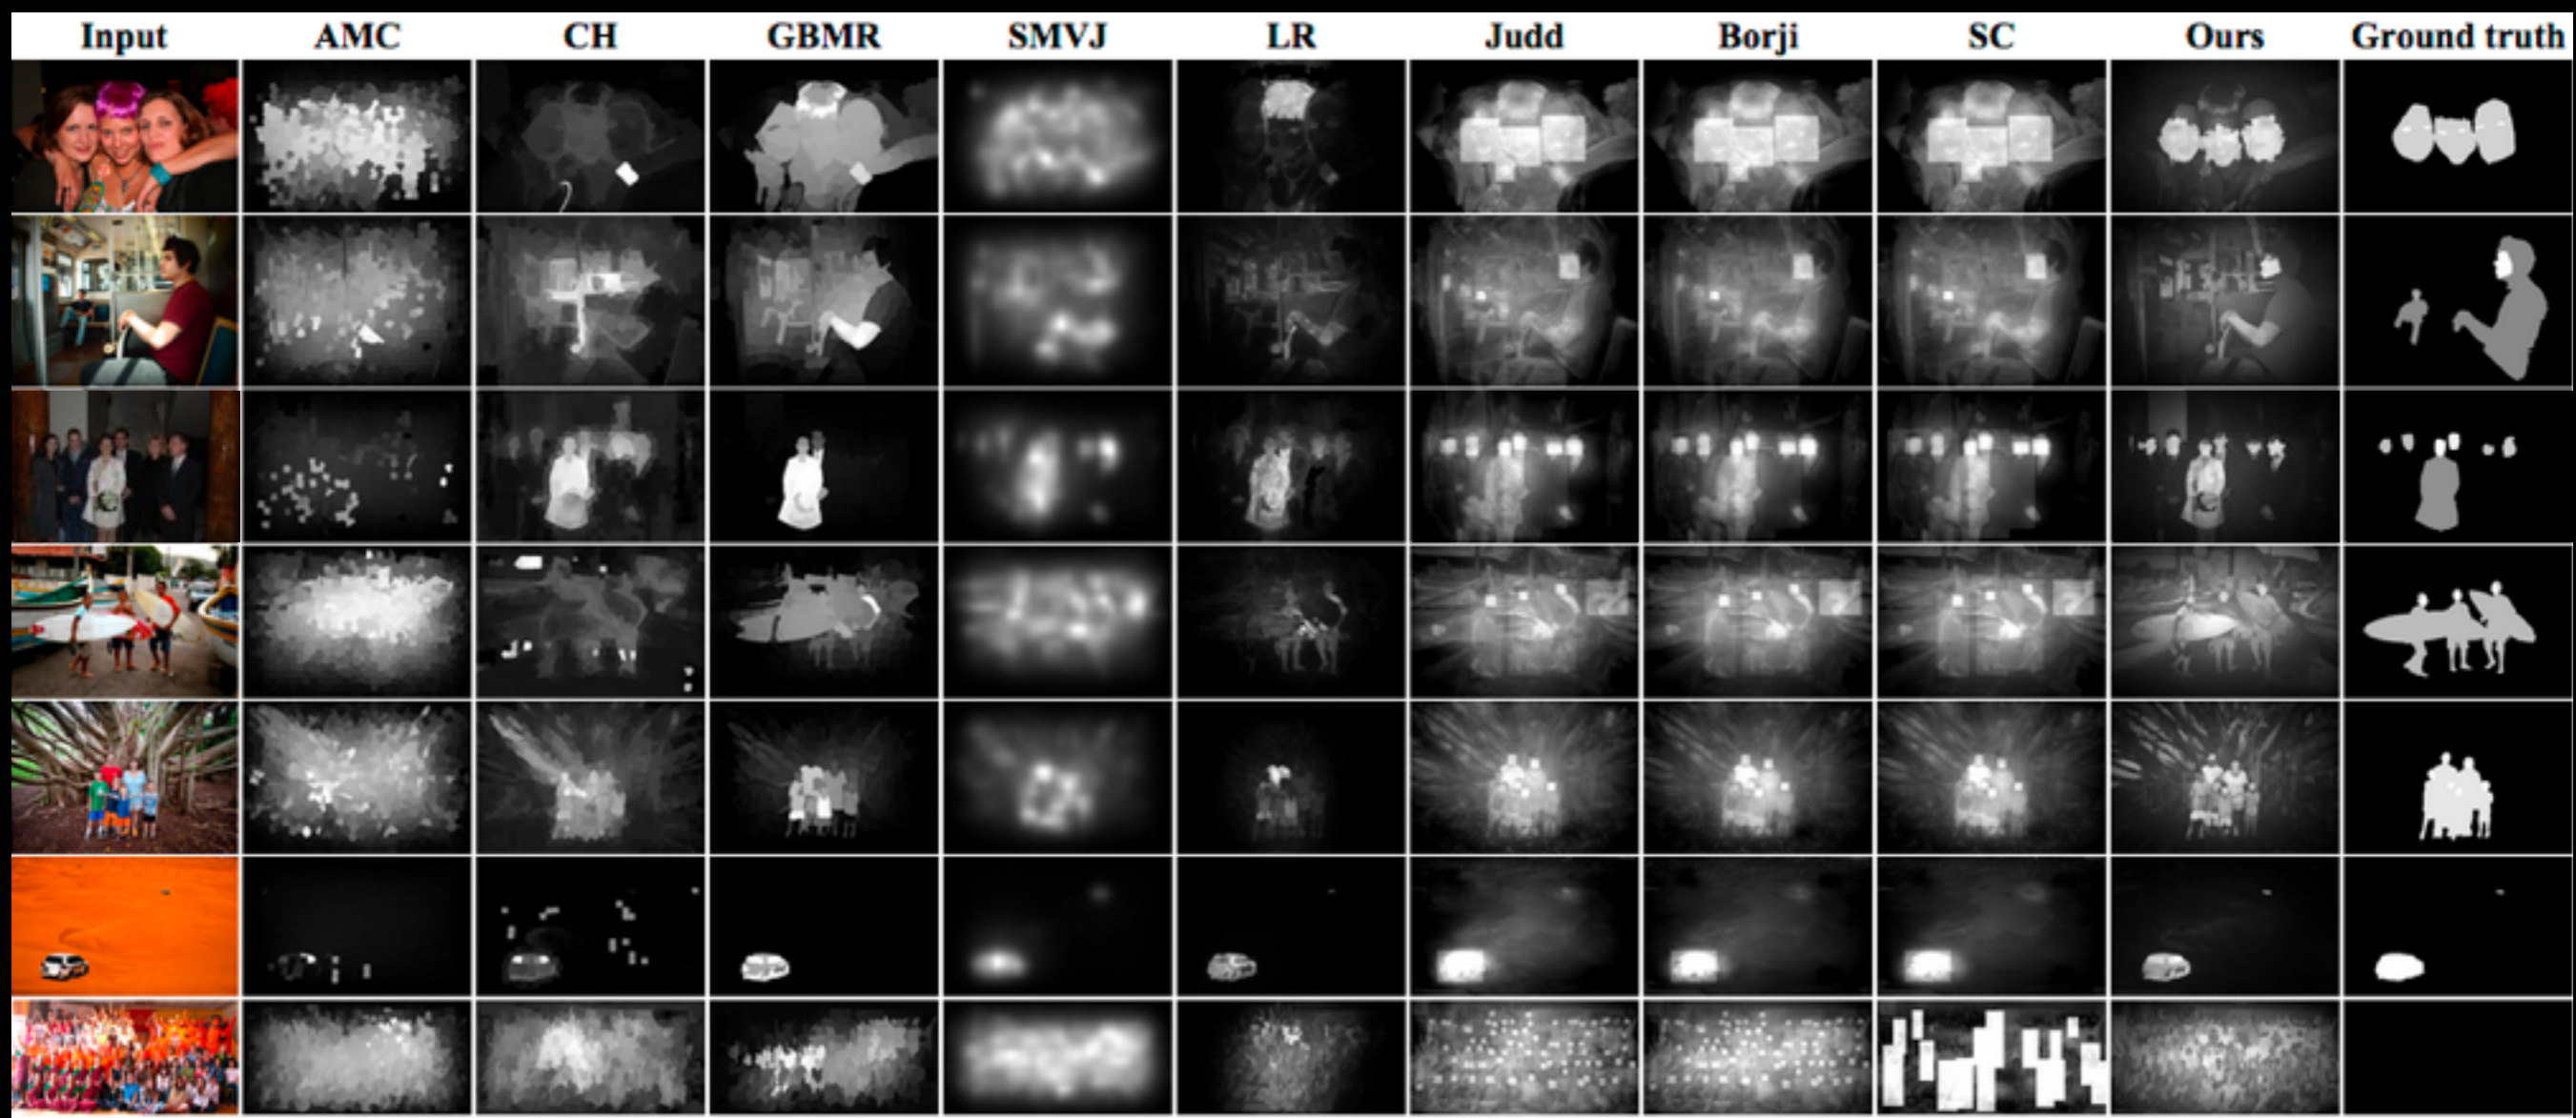
LR



Ground truth



Results




Results



LR	Judd	Borji	SC	Ours	Ground truth

Results




LR	Judd	Borji	SC	Ours	Ground truth

Results




LR



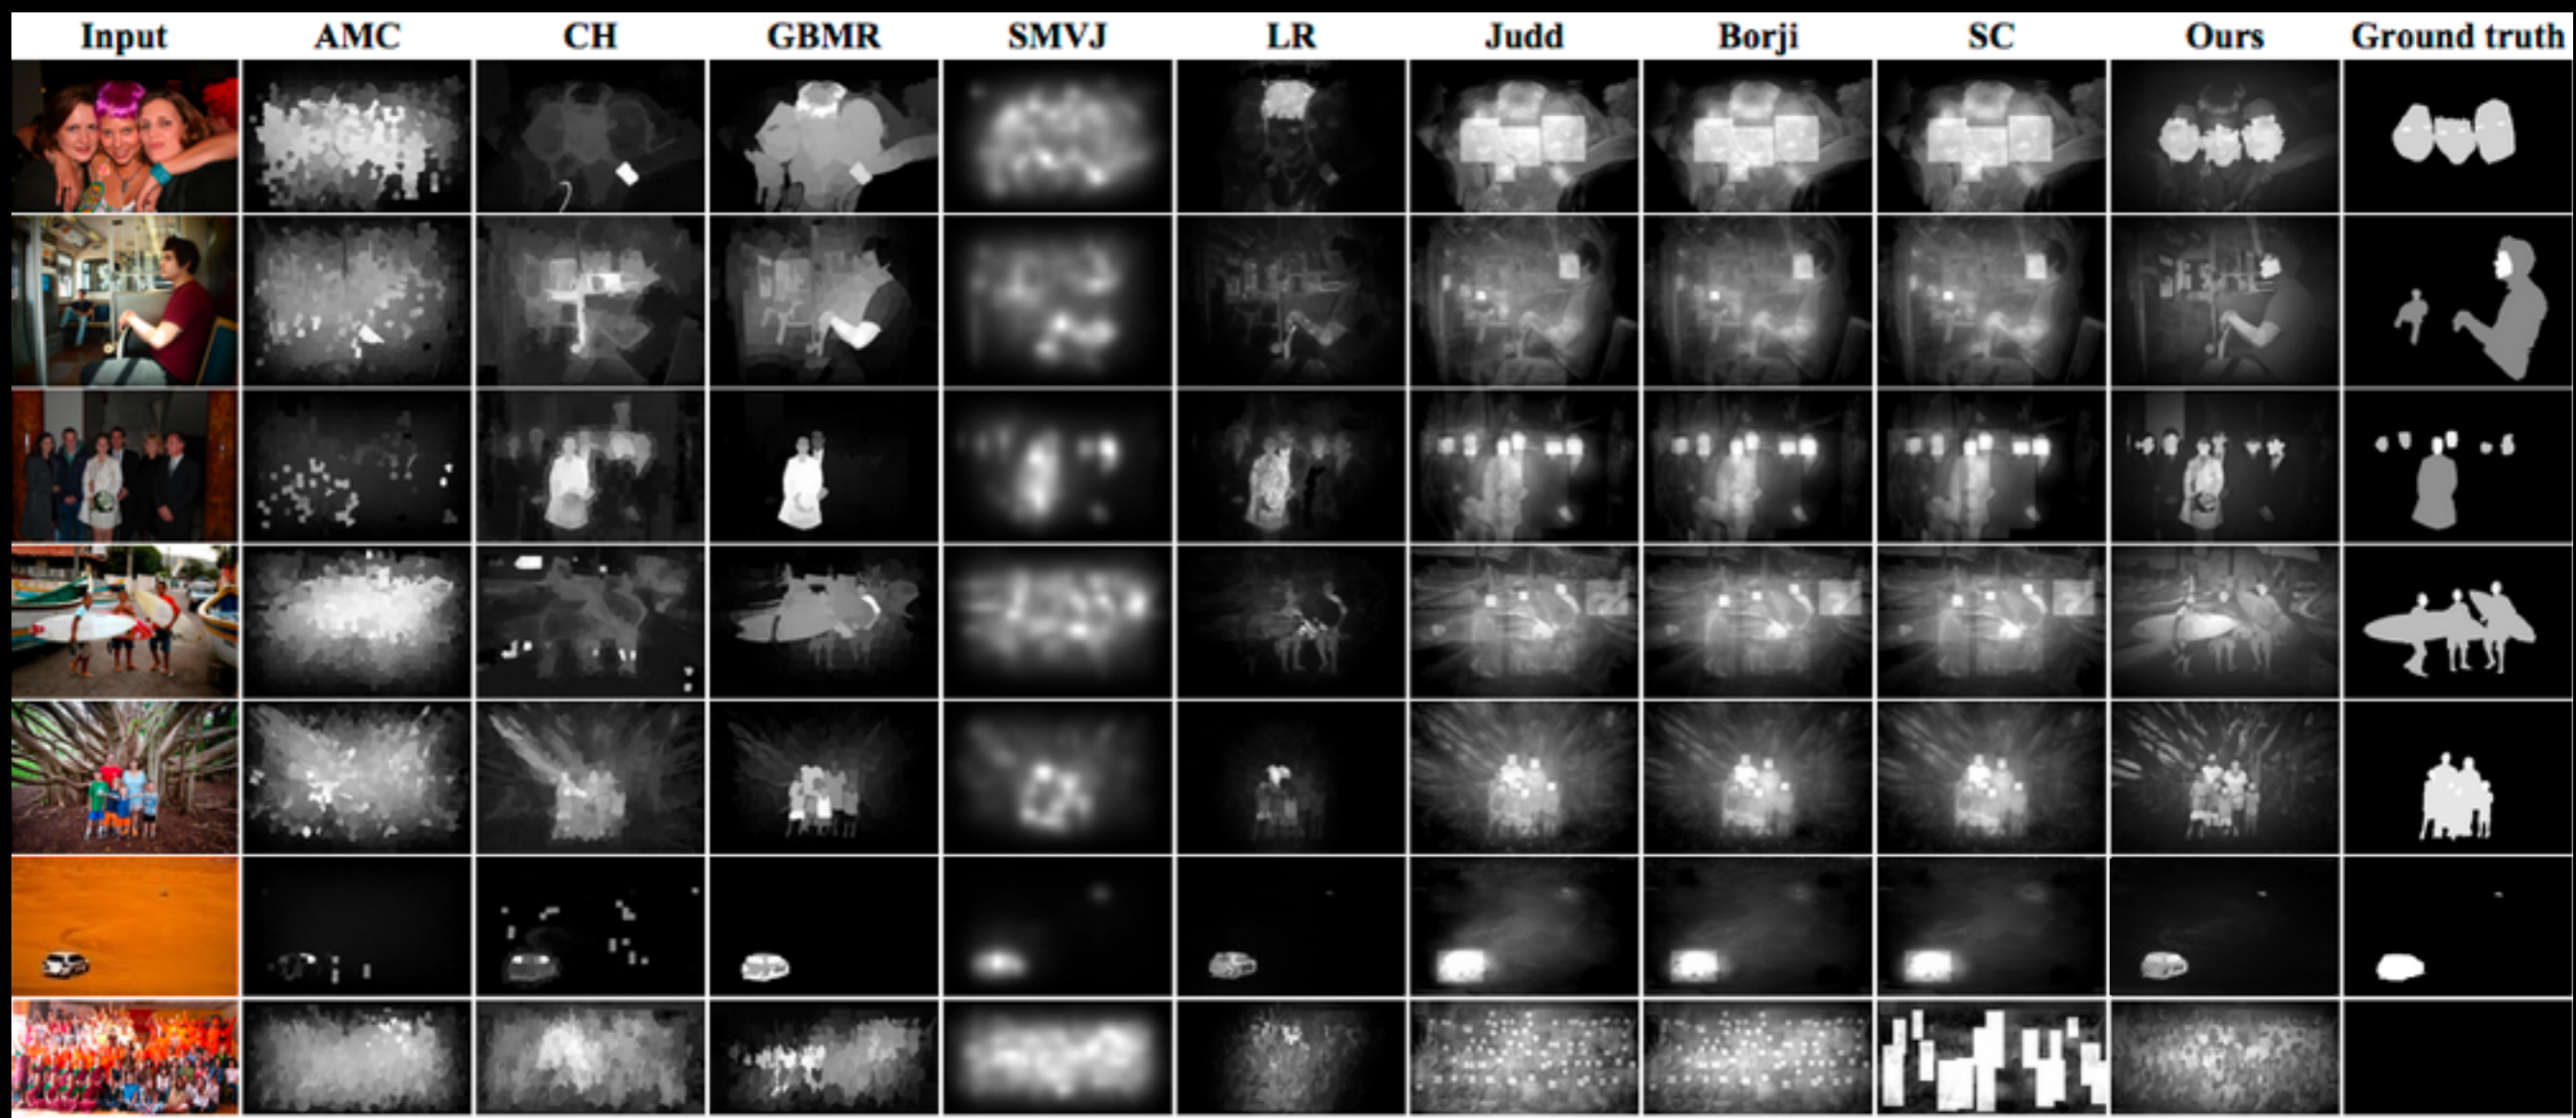
Results



LR



Results



Results



	LR	Judd	Borji	SC	Ours	Ground truth

Results




LR	Judd	Borji	SC	Ours	Ground truth

Results




LR



Results



LR



Conclusion


Conclusion




Conclusion



Conclusion



Conclusion




a face importance model

Conclusion



Conclusion




a face importance model




a saliency algorithm

Conclusion



Conclusion



a face importance model



a saliency algorithm



a saliency dataset

References

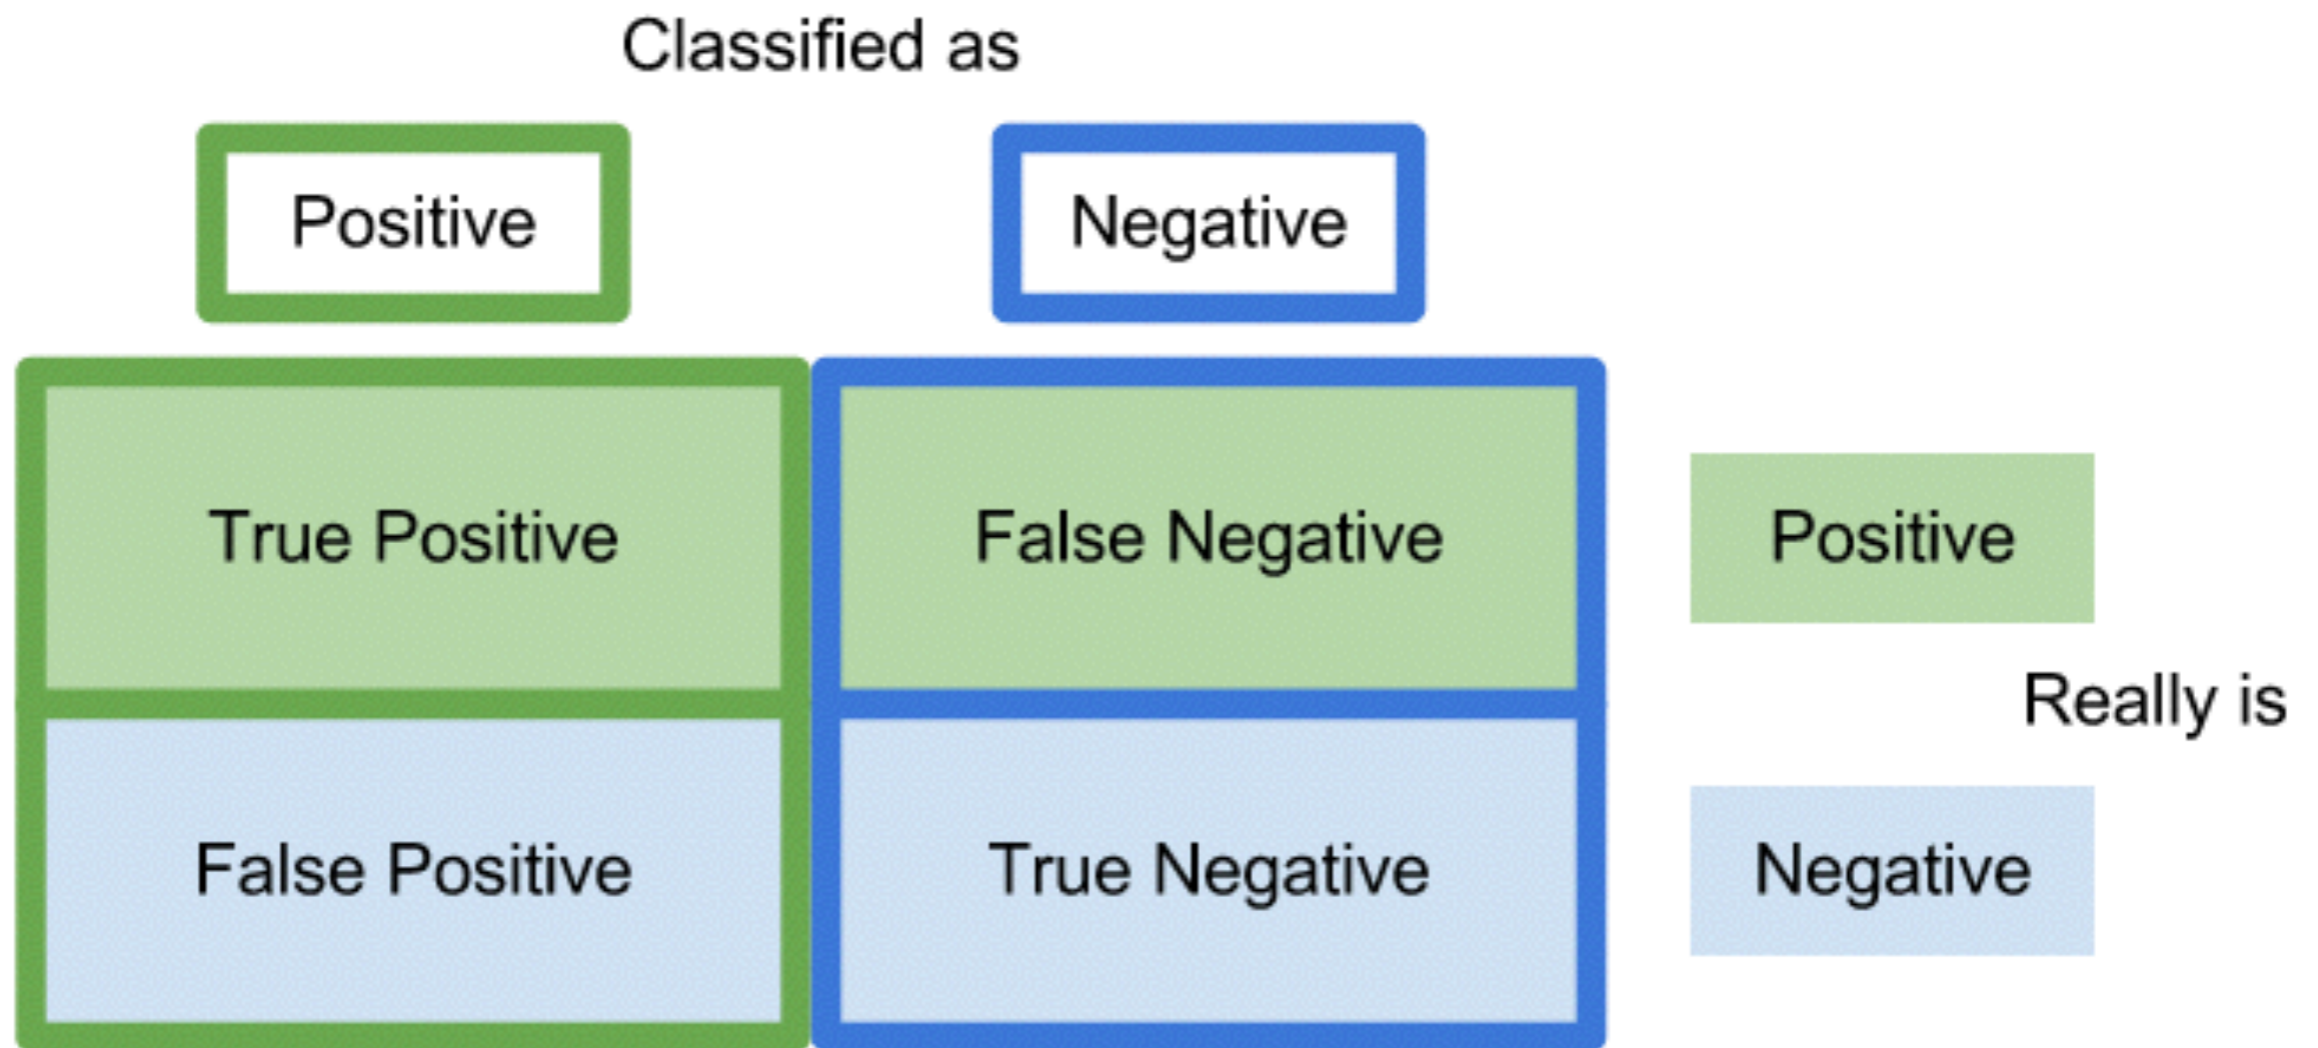
- SMVJ: Cerf, M., Harel, J., Einhäuser, W., and Koch, C., “Predicting human gaze using low-level saliency combined with face detection,” in [NIPS], (2007).
- Judd: Judd, T., Ehinger, K., Durand, F., and Torralba, A., “Learning to predict where humans look,” in [Proceedings of IEEE International Conference on Computer Vision], 2106–2113 (2009).
- AMC: Jiang, B., Zhang, L., Lu, H., Yang, M.-H., and Yang, C., “Saliency detection via absorbing markov chain,” Proceedings of IEEE International Conference on Computer Vision (2013).
- CH: Li, X., Li, Y., Shen, C., Dick, A., and van den Hengel, A., “Contextual hypergraph modelling for salient object detection,” Proceedings of IEEE International Conference on Computer Vision (2013).
- GBMR: Yang, C., Zhang, L., Lu, H., Ruan, X., and Yang, M.-H., “Saliency detection via graph-based manifold ranking,” in [Proceedings of IEEE Conference on Computer Vision and Pattern Recognition], 3166–3173 (2013).

References

- LR: Shen, X. and Wu, Y., “A unified approach to salient object detection via low rank matrix recovery,” in [Proceedings of IEEE Conference on Computer Vision and Pattern Recognition], 853–860 (2012).
- Borji: Borji, A., “Boosting bottom-up and top-down visual features for saliency estimation,” in [Proceedings of IEEE Conference on Computer Vision and Pattern Recognition], 438–445 (2012).
- SC: Karthikeyan, S., Jagadeesh, V., and Manjunath, B. S., “Learning top down scene context for visual attention modelling in natural images,” in [Proceedings of IEEE International Conference on Image Processing], (2013).
- Zhu, X. and Ramanan, D., “Face detection, pose estimation, and landmark localization in the wild,” Proceedings of IEEE Conference on Computer Vision and Pattern Recognition , 2879–2886 (2012).
- Viola, P. and Jones, M., “Rapid object detection using a boosted cascade of simple features,” Proceedings of IEEE Conference on Computer Vision and Pattern Recognition I, 511–518 (2001).

Thank you!

$$\text{precision} = \text{tp} / (\text{tp} + \text{fp})$$
$$\text{recall} = \text{tp} / (\text{tp} + \text{fn})$$



license

10014357906_9258a20b2f_o.jpg:
license: 4
url: http://farm6.staticflickr.com/5486/10014357906_9258a20b2f_o.jpg

10014783096_e5264e57c2_o.jpg:
license: 4
url: http://farm4.staticflickr.com/3724/10014783096_e5264e57c2_o.jpg

10057316635_23d3a8d982_o.jpg:
license: 4
url: http://farm4.staticflickr.com/3765/10057316635_23d3a8d982_o.jpg

10071994316_a9956df823_o.jpg:
license: 4
url: http://farm6.staticflickr.com/5501/10071994316_a9956df823_o.jpg

10100444523_2e76ab19f2_o.jpg:
license: 4
url: http://farm4.staticflickr.com/3810/10100444523_2e76ab19f2_o.jpg

10119159646_75c25d9520_o.jpg:
license: 5
url: http://farm6.staticflickr.com/5467/10119159646_75c25d9520_o.jpg

10119318984_dcb2df29ef_o.jpg:
license: 5
url: http://farm4.staticflickr.com/3685/10119318984_dcb2df29ef_o.jpg

10119466006_32b657f7ee_o.jpg:
license: 5
url: http://farm8.staticflickr.com/7304/10119466006_32b657f7ee_o.jpg

10313983086_81457222e5_o.jpg:
license: 4
url: http://farm4.staticflickr.com/3784/10313983086_81457222e5_o.jpg

10460818166_acec7a9d61_o.jpg:
license: 5
url: http://farm4.staticflickr.com/3809/10460818166_acec7a9d61_o.jpg

10460863345_26f834d8ac_o.jpg:
license: 5
url: http://farm6.staticflickr.com/5471/10460863345_26f834d8ac_o.jpg

license

- 0: All Rights Reserved
- 1: Attribution-NonCommercial-ShareAlike License
- 2: Attribution-NonCommercial License
- 3: Attribution-NonCommercial-NoDerivs License
- 4: Attribution License
- 5: Attribution-ShareAlike License
- 6: Attribution-NoDerivs License

Saliency vs importance

- Wang, J., Chandler, D. M., and Le Callet, P., “Quantifying the relationship between visual salience and visual importance,” IS&T/SPIE HVEI (2010): 80-90% overlap between importance and visual gaze.
- Crowd-sourcing needs to be easy.
- Application-oriented.