Beyond the Average Viewer

Predicting individual brain-response differences with TRIBE v2

Raphael Hauser[1], Young-Eun Lee[1], Chunyin Siu[1], Manish Saggar[1]

[1]Department of Psychiatry and Behavioral Sciences, Stanford University

Can we predict how one person's brain response to a movie differs from the group?

174 people watched the same short movie clips in an MRI scanner (Human Connectome Project data).

1 BACKGROUND

TRIBE v2 predicts brain responses to a movie from the movie itself.

- Predicting the average viewer well does not tell you how one person differs from it.
- Fine-tuning TRIBE v2 on someone's own scans does improve prediction of their whole response, but that accuracy is dominated by the shared part.
- For imaging biomarkers, the individual part has to be predictable at all first.

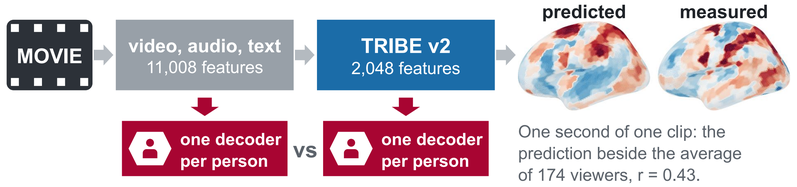

One second of one clip: the prediction beside the average of 174 viewers, r = 0.43.

We drop the model's own readout and fit our own decoder from each, on equal terms, with the same size and tuning grid for both.

frozen TRIBE v2 checkpoint: d'Ascoli et al., 2026 (arXiv:2605.04326)

2 IDEA

Everyone watching the same movie shares most of their brain response. Each person also deviates a little.

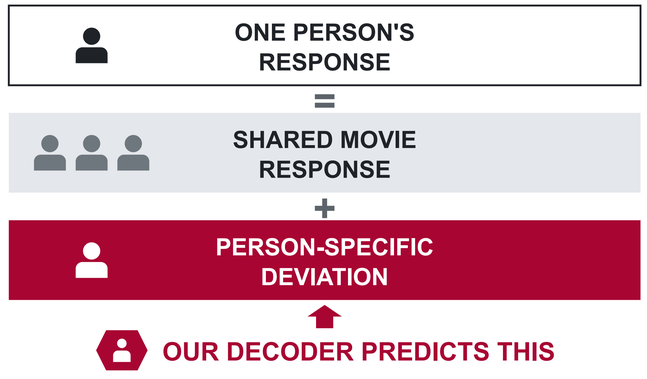

Intersubject component decomposition: the shared response is the average of participants outside your fold, relatives always held out together. Each person's share of it is fitted on their calibration runs; what is left is their deviation.

3 WHAT WE DID

Subtract the shared response, then learn each person's deviation from three runs and predict it on a fourth.

Our decoder is a map from the movie to one person: features go in, their deviation comes out at every cortical parcel.

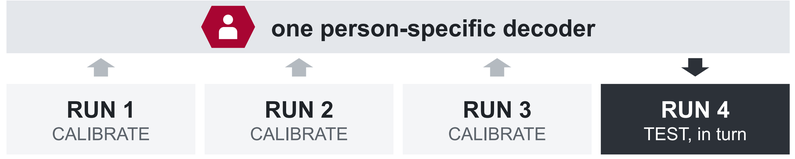

Each run takes its turn as the test run, never used for fitting. It holds different segments.

Each decoder is multi-output ridge on a low-rank projection of those features. The projection, the rank and the penalty are all fitted inside the calibration runs, on one grid for both representations.

174 participants · 89 family units · 4 HCP movie runs

Family units group relatives, so siblings never span a train/test split.
Schaefer-1000 parcels, 5 s response lag, family-aware 5-fold validation.

4 RESULT

On a run held out from fitting, your own decoder beats a stranger's, and TRIBE v2 strengthens the effect.

Participant specificity (r) = accuracy with your own decoder minus accuracy with a stranger's

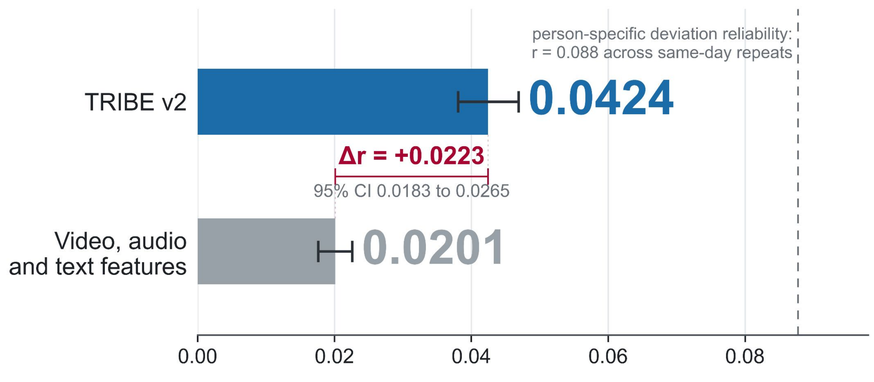

- ✔ Δr is positive in all 4 held-out runs
- ✔ Δr is positive in 80 of 89 family units
- ✔ holds when both are pinned to 10 components: Δr = +0.0260
- ✔ scrambling the movie's timing removes the effect (separate control, fixed decoder size)
- ✔ no other viewers' scans of the test run needed: Δr = +0.0215

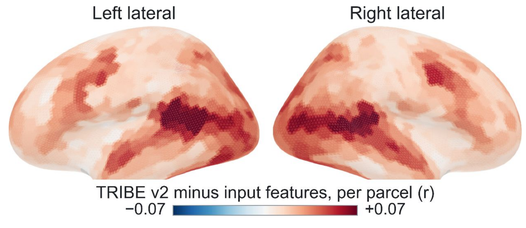

- Each of the 1000 parcels was scored on its own.
- Δr is positive in 991 of them, so no one region carries the effect.

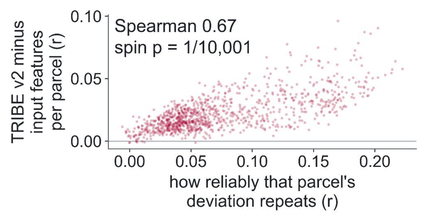

- The scatter plot puts each parcel's Δr against how reliably its deviation repeats, measured on a segment held out of the analysis.
- How strongly the movie drives a parcel tracks Δr as closely (0.69), raw response amplitude does not (0.41): the map follows measurability, not anatomy.

TRIBE v2 makes each person's deviation easier to read out.

Its 2,048 features are computed from the 11,008 it was given, so the gain is in the arrangement.

Every participant still needs their own movie fMRI to calibrate.

WHERE THIS COULD LEAD

- ➡ A model of your cortex that generalizes to films it has never seen.
- ➡ Minutes of film where a personalized map now takes hours of task fMRI.
- ➡ The variation that group averaging throws away, treated as the signal.
- ➡ Individual differences reliable enough to be worth testing in the clinic.

raphaelhauser.com/tribe

poster PDF · code and preprint when published