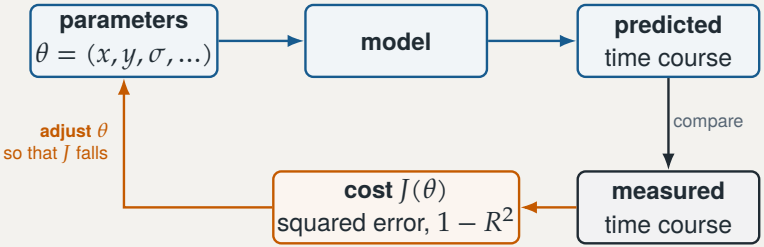
POPYLAR PRF WORKSHOP · COIMBRA · FRIDAY 6 NOVEMBER 2026

ADVANCED PRF MODELS

Where, what, how and when: from population receptive fields to richer encoding models

Tomas Knapen

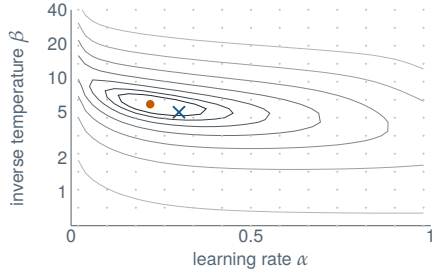
ABOUT FITTING MODELS



every fitter in prfmodel, grid search, least squares or gradient descent, runs this loop

- Before the richer models: what a fit is, and what can go wrong.

SEARCH EVERY POINT OF A GRID



negative log likelihood of 200 choices

contours at 1, 2, 4, …, 64 above the minimum

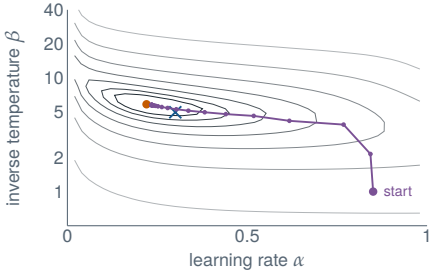
× the generating values, $\alpha = 0.3, \beta = 5$

• the grid minimum: $\alpha = 0.22, \beta = 5.9$

dots: every 4th grid point of 2401

- Grid search: evaluate the cost at every combination, keep the lowest.
- Two parameters at 49 values: 2 401 evaluations; three: 117 649; eight: an explosion.

FOLLOW THE SLOPE DOWN



the same landscape, walked

contours at 1, 2, 4, …, 64 above the minimum

× the generating values, $\alpha = 0.3$, $\beta = 5$

• the grid minimum: $\alpha = 0.22$, $\beta = 5.9$

gradient descent, 25 steps from a poor start

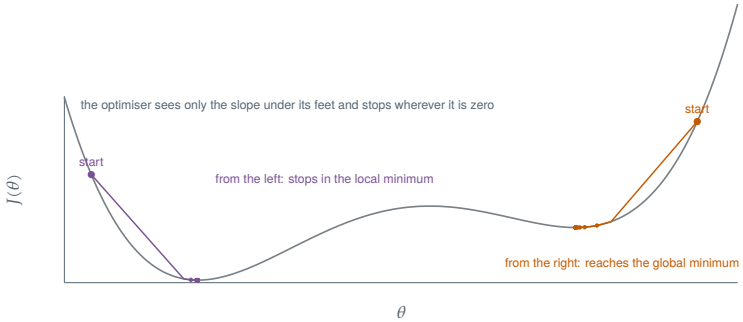
- Gradient descent: start anywhere, measure the slope, step against it, repeat; 25 steps where the grid used 169.
- The step is the slope times a learning rate; too large a rate overshoots.

THE STEP SIZE IS A CHOICE



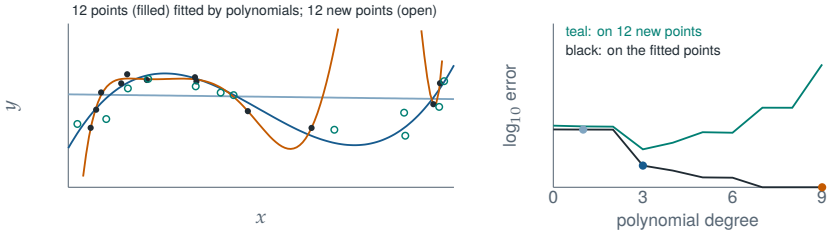
- Set up a cost $J(\theta)$, and on each step update $\theta \leftarrow \theta - \eta \, \mathrm{d}J/\mathrm{d}\theta$.
- Too small: many steps. Too large: overshoots, can diverge. About right: a few steps.

STUCK IN A LOCAL MINIMUM



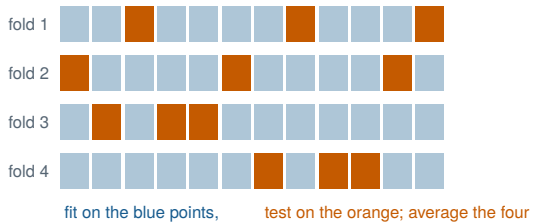
- Descent follows the slope under its feet and stops wherever it is flat.
- No general cure: start several times, keep the best, check the result.

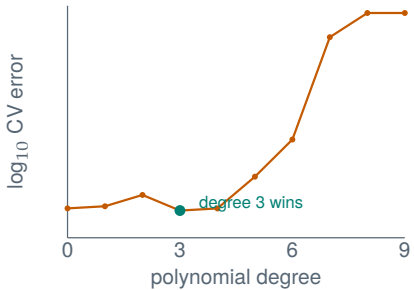
FITTING THE NOISE



- Degree 9 passes through every fitted point and misses the new ones: it learned the noise.
- Degree 3 does best on new points, and only that error counts.

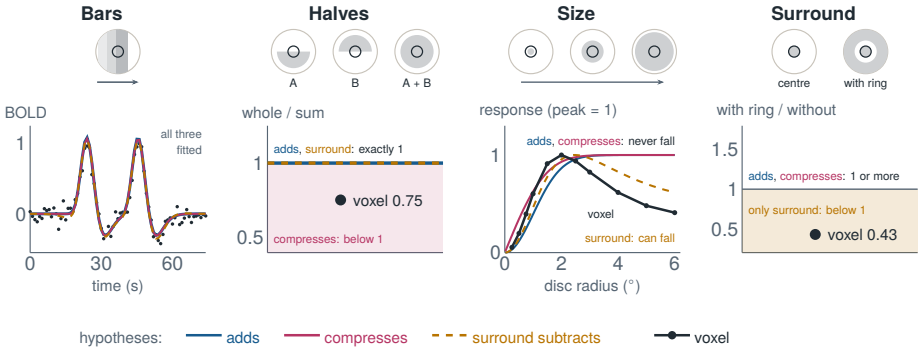
CROSS-VALIDATION





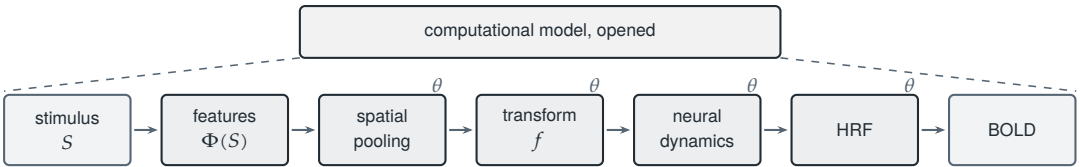
- Fit on the training folds, score on the test fold; rotate, average.
- Noise fitted on training fails on the test fold: degree 3 wins here too.

EXPLAINING YESTERDAY'S VOXEL



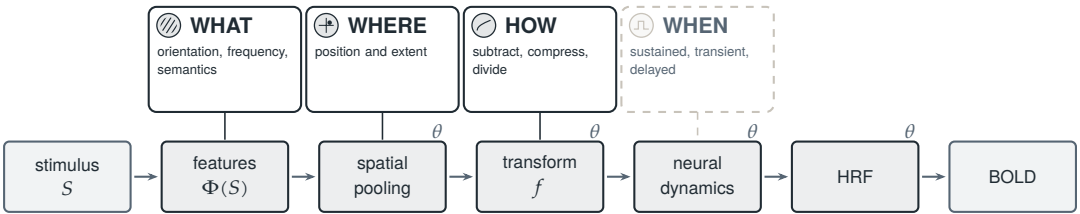
- We designed the experiment; now we have to explain the response.

EVERY PRF IS AN ENCODING MODEL



- A pRF model is a map from stimulus to predicted measurement, with parameters to fit and to test.

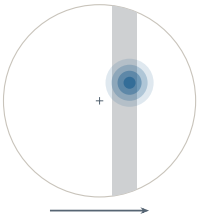
WHERE, WHAT, HOW, WHEN

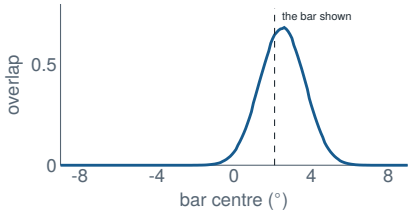


Four questions about one model:

- where it pools
- what it is tuned to
- how it transforms
- when it integrates

A GAUSSIAN SUMS OVER SPACE

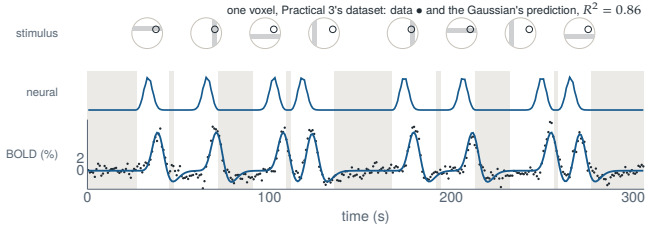




$$r(t) = \sum_{x,y} G(x, y;\, x_0, y_0, \sigma)\, S(x, y, t)$$

- The response is a weighted sum of stimulation over space; the weights are a Gaussian. The bar's aperture is the stimulus S: one where it is, zero elsewhere.

FROM STIMULUS TO BOLD



$$r(t) = \sum_{x,y} G(x, y;\, x_0, y_0, \sigma)\, S(x, y, t), \qquad \mathrm{BOLD}(t) = (h * r)(t)$$

- One set of parameters predicts the whole time course; that prediction is what we fit.

WHAT THE GAUSSIAN FORBIDS

Parts add



$r(A+B) = r(A) + r(B)$

More is never less



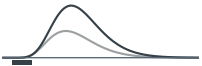
every weight positive:
added stimulation adds

Only where counts



same aperture,
same response

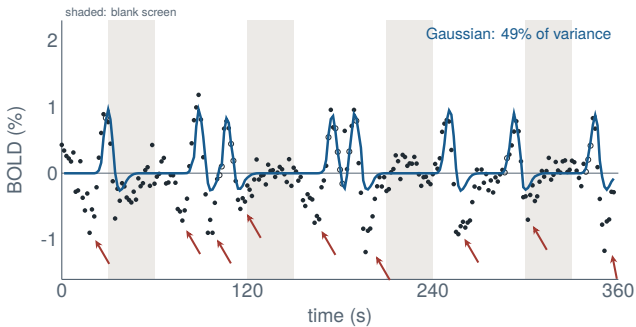
Activations sum in time

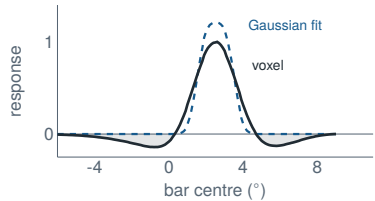


twice as long,
twice the response

- All models are wrong, but some are useful *(Box 1979)*: useful because they forbid things; each forbidden response is a test.

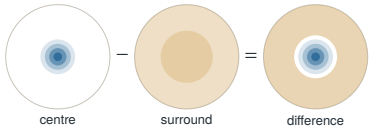
MORE STIMULUS, LESS RESPONSE





- A bar beside the centre can push the response below baseline; no positive-weighted sum does that *(Zuiderbaan et al. 2012)*.

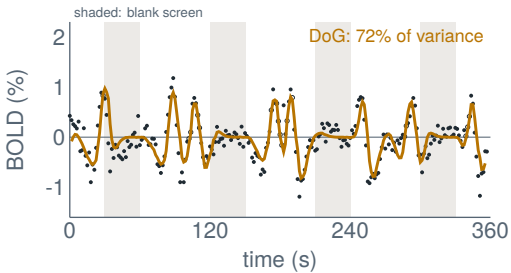
CENTRE MINUS SURROUND

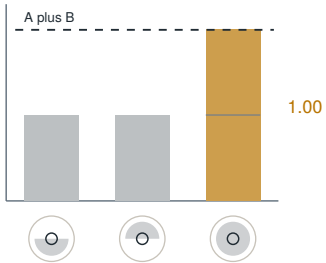


$$r(t) = \sum_{x,y} \left[a_c \, G_c(x,y) - a_s \, G_s(x,y) \right] S(x,y,t)$$

- DoG subtracts a broader Gaussian: stimulation around the centre now lowers the response.

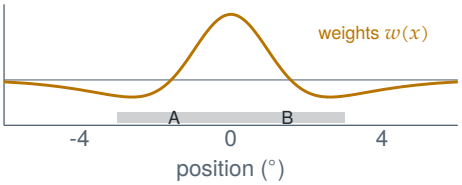
A SURROUND, BUT STILL A SUM





- DoG captures the negative lobe but still adds: the whole equals the sum of its parts.

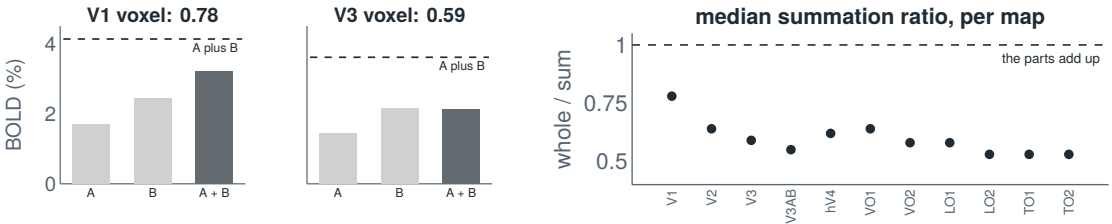
ANY LINEAR FILTER ADDS



$$r(A) = \sum_x w(x)\, A(x)$$
$$r(B) = \sum_x w(x)\, B(x)$$
$$r(A{+}B) = \sum_x w(x)\, [A(x) + B(x)]$$
$$= r(A) + r(B)$$

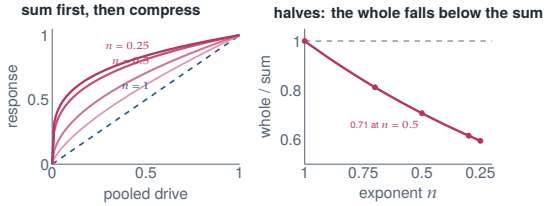
- For any weights w, positive or negative:
  $r(A + B) = \sum w\,(A + B) = \sum wA + \sum wB = r(A) + r(B).$

THE WHOLE IS LESS THAN THE SUM



- Yesterday's experiment is today's constraint: no linear model falls below the sum *(Kay et al. 2013)*.

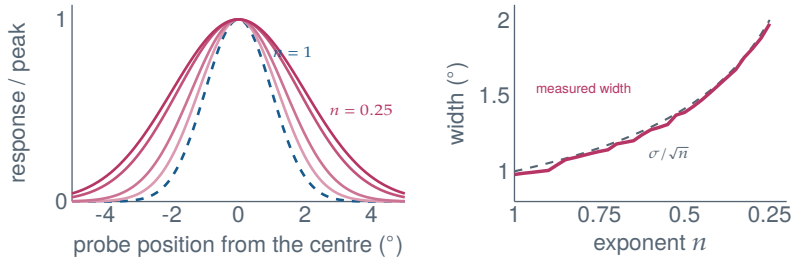
SUM FIRST, THEN COMPRESS



$$r(t) = \Big[\sum_{x,y} G(x,y)\, S(x,y,t) \Big]^{n}, \qquad 0 < n \leq 1$$

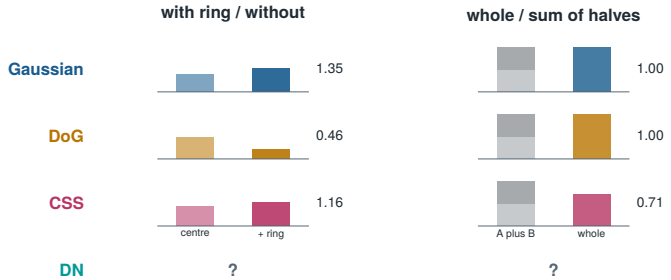
- CSS raises the summed drive to a power below one; the whole then falls below the sum.

COMPRESSION WIDENS THE FIELD



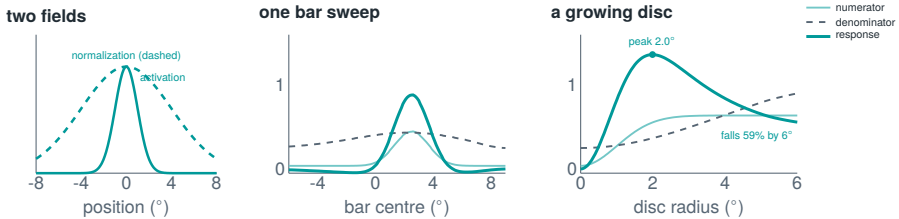
- With $n < 1$ the response to a small probe falls off more slowly than the Gaussian, as G^n: the effective size is $\sigma / \sqrt{n}$, larger than σ *(Kay et al. 2013)*.

TWO FAILURES, TWO FIXES



- DoG gives the negative lobe, CSS the compression; neither gives both.

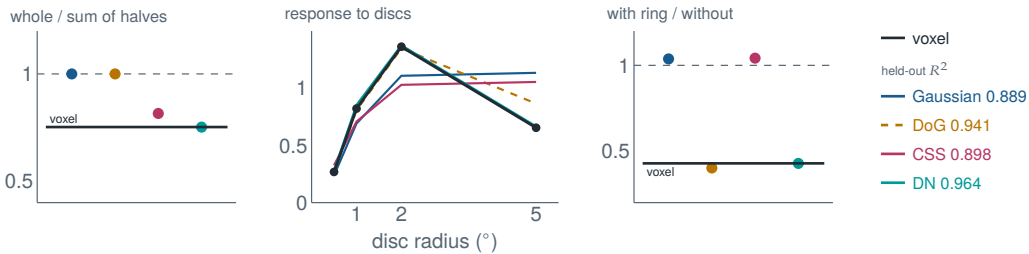
DIVIDE BY THE CONTEXT



$$r = \frac{a\,A(S) + b}{c\,N(S) + d} - \frac{b}{d}$$

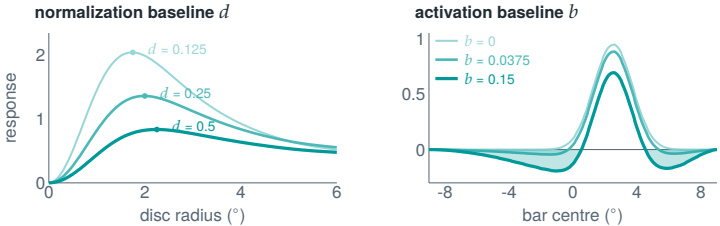
- A canonical computation *(Carandini and Heeger 2012)*: a small field's drive over a larger field's.

ONE OPERATION, BOTH SIGNATURES



- One divisive operation fits suppression and compression in the same voxel, as in measured voxels *(Aqil et al. 2021)*.

PARAMETERS SHAPE RESPONSES



$$r = \frac{a\,A(S) + b}{c\,N(S) + d} - \frac{b}{d}$$

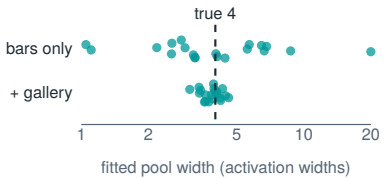
- d sets the oversaturation and the compression; b sets the surround.

PINNING THE PARAMETERS

Scaling a, b, c and d by any factor k leaves the prediction unchanged.

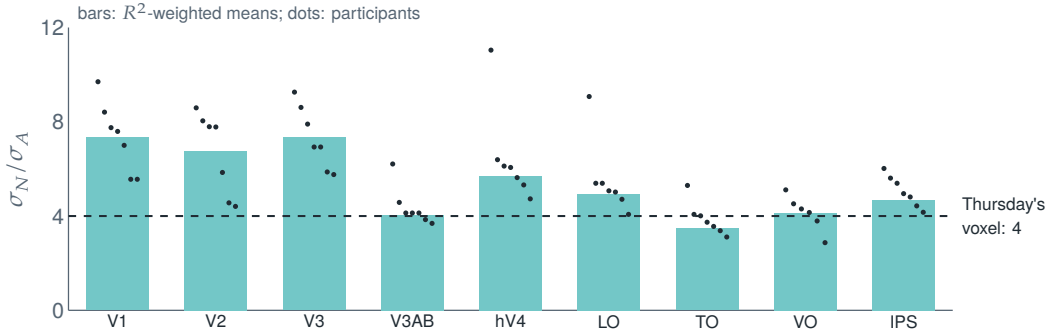
$$\frac{ka\,A + kb}{kc\,N + kd} - \frac{kb}{kd} = \frac{aA + b}{cN + d} - \frac{b}{d}$$

any factor k: same prediction



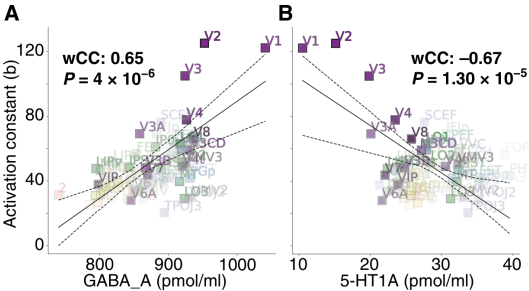
- Fix one scale ($c = 1$); probe the pool with stimuli that reach it.

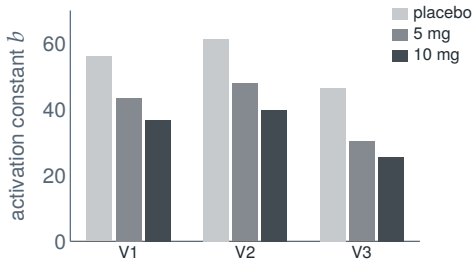
ACROSS THE VISUAL HIERARCHY



- The normalization pool's width changes systematically across visual areas *(Aqil et al. 2021)*.

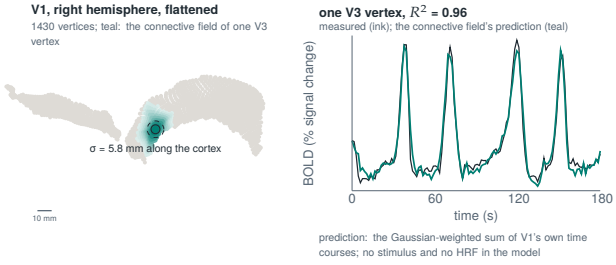
FROM PARAMETERS TO BIOLOGY





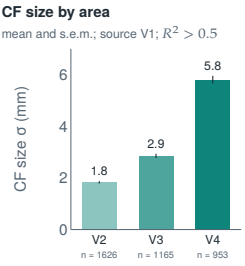
- b tracks $GABA_A$ and $5\text{-}HT_{1A}$ receptor densities across areas *(Aqil et al. 2024)*; psilocybin lowers it *(Aqil et al. 2025)*.

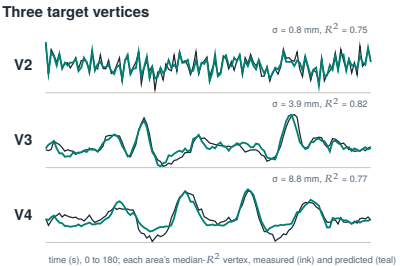
CONNECTIVE FIELDS: PRFS ON THE CORTEX



- A Gaussian over cortical distance, fed by V1's activity *(Haak et al. 2013)*.

FIELDS GROW UP THE HIERARCHY

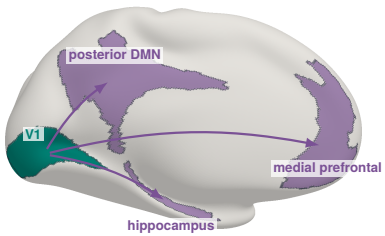




- 1.8 mm in V2, 2.9 in V3, 5.8 in V4: each area pools a wider patch of V1.

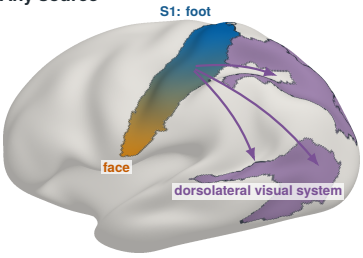
ANY STIMULUS, ANY SOURCE

Any stimulus



Knapen (2021): V1 as the source during movie, rest and mapping

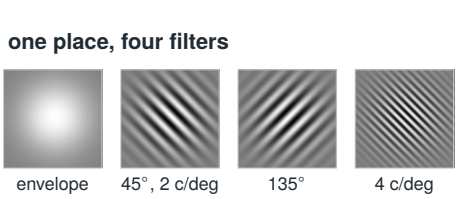
Any source

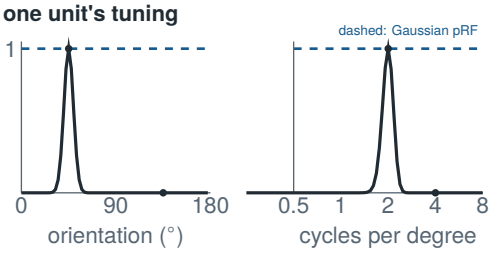


Hedger et al. (2025): the body map of S1 as the source during video watching

- The stimulus is the source's activity: movie, rest and mapping all work *(Knapen 2021)*.
- The source can be any map: body tuning from S1 tiles vision *(Hedger et al. 2025)*.

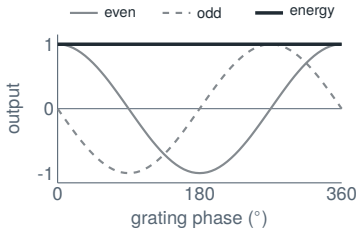
SAME PLACE, OTHER FEATURES

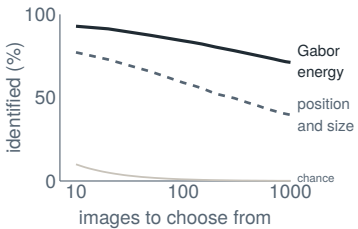




- Two images at one place can drive a voxel differently: position is not the whole stimulus.

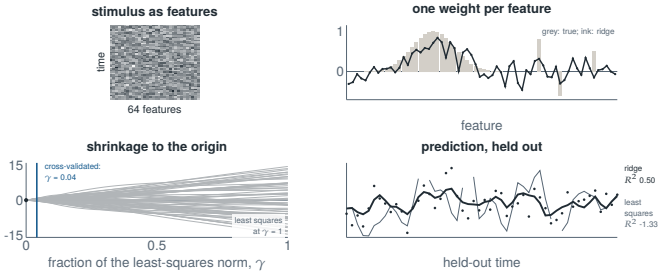
A RECEPTIVE FIELD FOR IMAGES





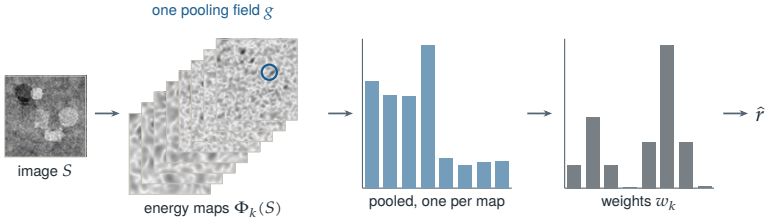
- Two Gabor filters in quadrature, squared and summed, respond to contrast whatever its phase. Gabor energy pooled over space predicted responses to new photographs, well enough to identify them: 92% and 72% of 120 new images in two subjects, where chance is 0.8% *(Kay et al. 2008)*.

THOUSANDS OF FEATURES



- The weights become the description *(Huth et al. 2016)*; ridge shrinks them toward the origin *(Rokem and Kay 2020)*.
- A different approach from the few-parameter process models in prfmodel.

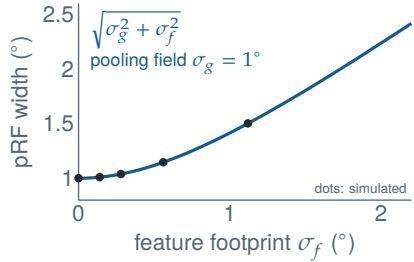
THE FEATURE-WEIGHTED RF



$$\hat{r} = \sum_k w_k \sum_{x,y} g(x, y;\, \theta)\, \Phi_k(S, x, y)$$

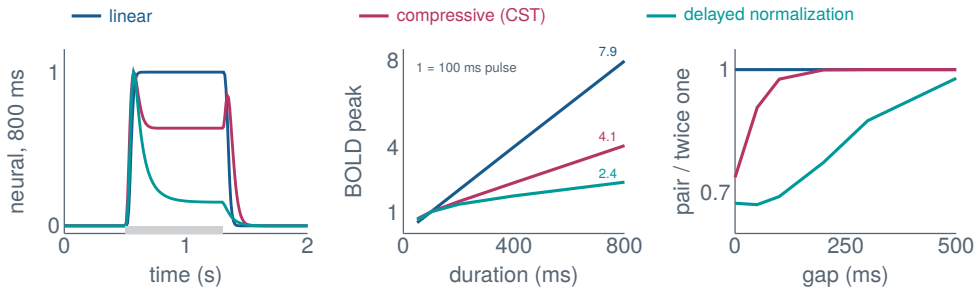
- One shared pooling field, features as rich as needed: where times what *(St-Yves and Naselaris 2018)*.

POOLING WIDTH IN VISUAL SPACE



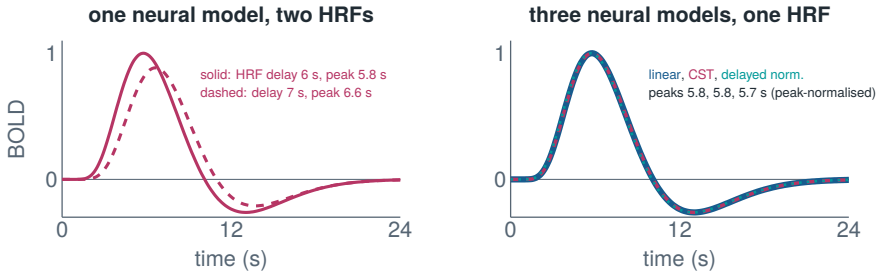
- The field pools feature maps whose filters have their own footprint σ_f: in visual degrees the pRF is $\sqrt{\sigma_g^2 + \sigma_f^2}$, so the pooling field alone understates its size *(St-Yves and Naselaris 2018)*.

SAME PLACE, OTHER TIMING



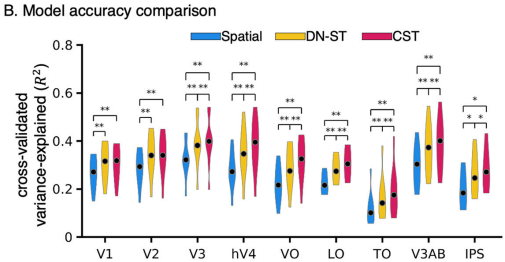
- Same place, same features, different timing: space alone does not set the response.

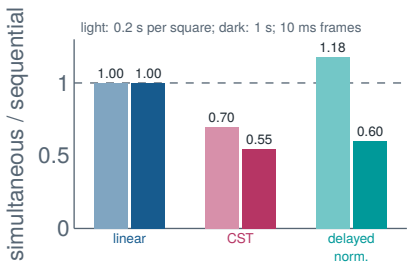
THE HRF IS NOT NEURAL TIME



- A slower HRF moves the BOLD peak for every stimulus alike (5.8 to 6.6 s here); the neural models, through one HRF, give the same BOLD shape and differ in how the amplitude grows with duration, which a linear HRF cannot mimic.

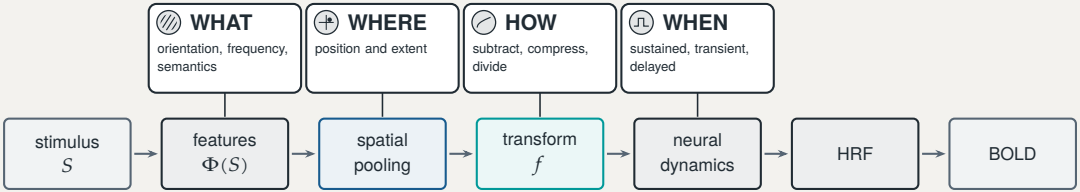
SPATIOTEMPORAL PRFS





- Spatiotemporal pRF models exist and fit fMRI *(Kim et al. 2024; Kupers et al. 2024)*; prfmodel ships two in its beta release.

MODELS FOR THE QUESTION



1. Specify a predictive hypothesis.
2. Design stimuli that constrain its parameters.
3. Compare predictions on data not used to fit it.

Practical 3, today at 12:00

Gaussian — DoG — CSS — DN

- prfmodel implements these models and their components; time to play with them in practice.

REFERENCES I

Aqil, Marco, Gilles de Hollander, Nina Vreugdenhil, Tomas Knapen, and Serge O. Dumoulin. 2025. “Psilocybin Alters Visual Contextual Computations.” *Nature Communications* 16 (1): 10271. https://doi.org/10.1038/s41467-025-65150-y.

Aqil, Marco, Tomas Knapen, and Serge O. Dumoulin. 2021. “Divisive Normalization Unifies Disparate Response Signatures Throughout the Human Visual Hierarchy.” *Proceedings of the National Academy of Sciences* 118 (46): e2108713118. https://doi.org/10.1073/pnas.2108713118.

Aqil, Marco, Tomas Knapen, and Serge O. Dumoulin. 2024. “Computational Model Links Normalization to Chemoarchitecture in the Human Visual System.” *Science Advances* 10 (1): eadj6102. https://doi.org/10.1126/sciadv.adj6102.

REFERENCES II

Box, George E. P. 1979. “Robustness in the Strategy of Scientific Model Building.” In *Robustness in Statistics*, edited by Robert L. Launer and Graham N. Wilkinson. Academic Press. https://doi.org/10.1016/B978-0-12-438150-6.50018-2.

Carandini, Matteo, and David J. Heeger. 2012. “Normalization as a Canonical Neural Computation.” *Nature Reviews Neuroscience* 13 (1): 51–62. https://doi.org/10.1038/nrn3136.

Haak, Koen V., Jonathan Winawer, Ben M. Harvey, et al. 2013. “Connective Field Modeling.” *NeuroImage* 66: 376–84. https://doi.org/10.1016/j.neuroimage.2012.10.037.

Hedger, Nicholas, Thomas Naselaris, Kendrick Kay, and Tomas Knapen. 2025. “Vicarious Body Maps Bridge Vision and Touch in the Human Brain.” *Nature* 650 (8100): 173–81. https://doi.org/10.1038/s41586-025-09796-0.

REFERENCES III

Huth, Alexander G., Wendy A. de Heer, Thomas L. Griffiths, Frédéric E. Theunissen, and Jack L. Gallant. 2016. “Natural Speech Reveals the Semantic Maps That Tile Human Cerebral Cortex.” *Nature* 532 (7600): 453–58. https://doi.org/10.1038/nature17637.

Kay, Kendrick N., Thomas Naselaris, Ryan J. Prenger, and Jack L. Gallant. 2008. “Identifying Natural Images from Human Brain Activity.” *Nature* 452 (7185): 352–55. https://doi.org/10.1038/nature06713.

Kay, Kendrick N., Jonathan Winawer, Aviv Mezer, and Brian A. Wandell. 2013. “Compressive Spatial Summation in Human Visual Cortex.” *Journal of Neurophysiology* 110 (2): 481–94. https://doi.org/10.1152/jn.00105.2013.

Kim, Insub, Eline R. Kupers, Garikoitz Lerma-Usabiaga, and Kalanit Grill-Spector. 2024. “Characterizing Spatiotemporal Population Receptive Fields in Human Visual Cortex with fMRI.” *Journal of Neuroscience* 44 (2): e0803232023. https://doi.org/10.1523/JNEUROSCI.0803-23.2023.

REFERENCES IV

Knapen, Tomas. 2021. “Topographic Connectivity Reveals Task-Dependent Retinotopic Processing Throughout the Human Brain.” *Proceedings of the National Academy of Sciences* 118 (2): e2017032118. https://doi.org/10.1073/pnas.2017032118.

Kupers, Eline R., Insub Kim, and Kalanit Grill-Spector. 2024. “Rethinking Simultaneous Suppression in Visual Cortex via Compressive Spatiotemporal Population Receptive Fields.” *Nature Communications* 15 (1): 6885. https://doi.org/10.1038/s41467-024-51243-7.

Rokem, Ariel, and Kendrick Kay. 2020. “Fractional Ridge Regression: A Fast, Interpretable Reparameterization of Ridge Regression.” *GigaScience* 9 (12): giaa133. https://doi.org/10.1093/gigascience/giaa133.

REFERENCES V

St-Yves, Ghislain, and Thomas Naselaris. 2018. “The Feature-Weighted Receptive Field: An Interpretable Encoding Model for Complex Feature Spaces.” *NeuroImage* 180: 188–202. https://doi.org/10.1016/j.neuroimage.2017.06.035.

Zuiderbaan, Wietske, Ben M. Harvey, and Serge O. Dumoulin. 2012. “Modeling Center-Surround Configurations in Population Receptive Fields Using fMRI.” *Journal of Vision* 12 (3): 10. https://doi.org/10.1167/12.3.10.