Human visual object categorization is best described by a model with few stored exemplars

Robert J. Peters (1), Alex Backer (2), Fabrizio Gabbiani (3), and Christof Koch (1)
(1) Computation and Neural Systems, (2) Biology, Caltech, Pasadena, CA 91125
(3) Baylor College of Medicine, Houston, TX 77030

1 Introduction

- Models of visual object categorization:
  - ***all-exemplar models*** (e.g., GCM; Nosofsky, 1991)
  - ***prototype models*** (e.g., Reed, 1972)
  - ***decision boundary models*** (e.g., Maddox & Ashby, 1993)
- Very often the all-exemplar models win out—why?
  - *memory capacity?*
  - *orientation of decision boundary?*
  - *shape of decision boundary?*
- Past comparisons of models have not resolved these factors
- We introduced the ***roaming exemplar model*** to help sort them out

2 Brunswik faces

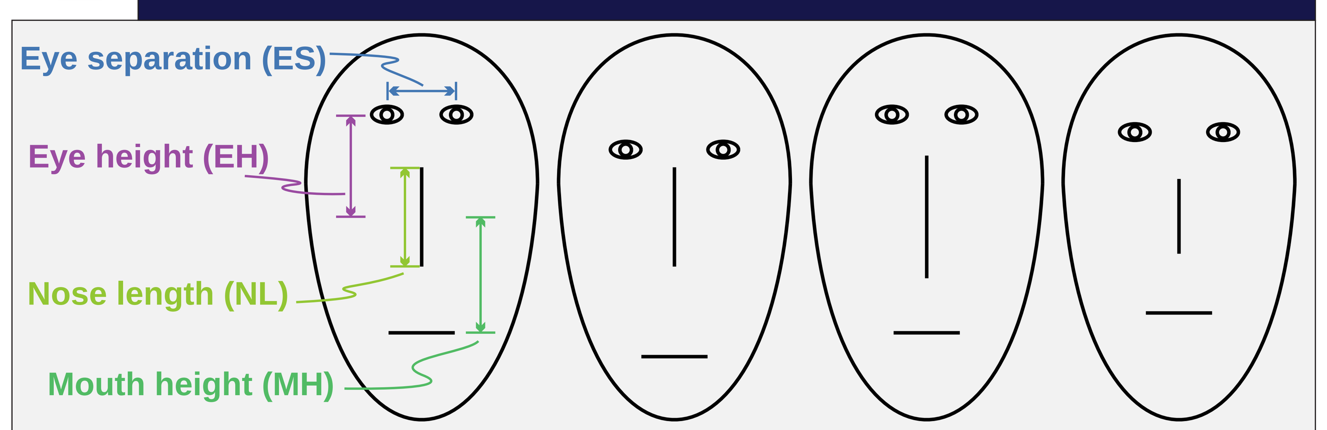

3 Categorization task

❶ Training phase
- uses category 1 & 2 training exemplars
- presented one at a time in random order
- subject guesses category 1 or 2
- auditory feedback is given
- repeat until subject reaches 85% correct

❷ Testing phase
like testing phase, except:
- also uses additional test exemplars
- feedback is not given
- repeat until each exemplar is seen 7 times
- subjects' responses are used to fit the models

9 subjects did training and testing for each of 12 sets of categories

4 Twelve sets of categories

- 10 training exemplars in category 1
- 10 training exemplars in category 2
- 60 test exemplars (not shown)
- 12 sets of categories, with different boundaries

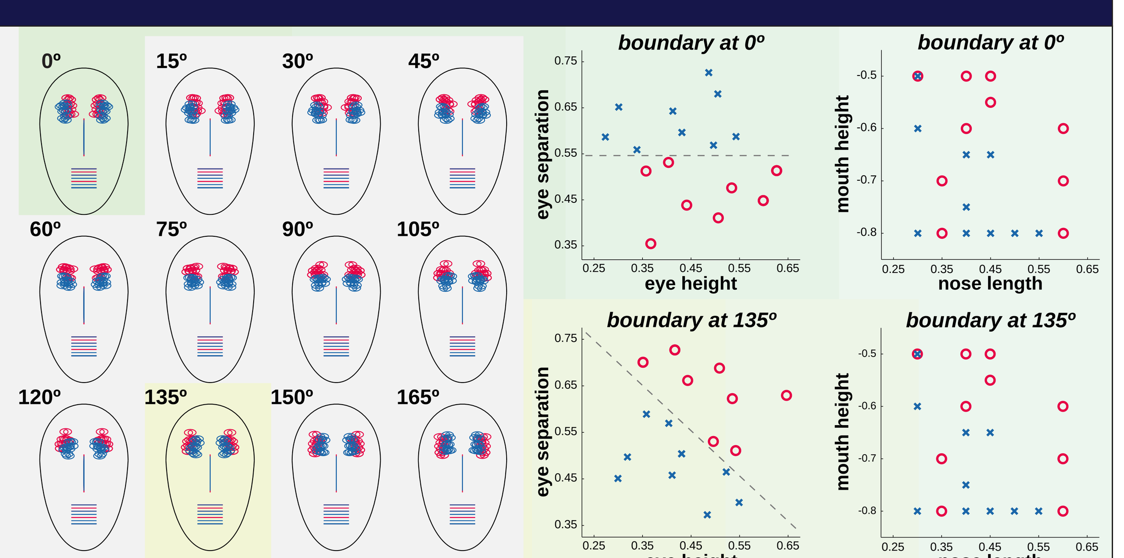

5 Categorization models

The RXM[1] has the set of properties that allow it to best fit human behavior

Boundary models
Categories are represented as Gaussian distributions, so decision boundaries are linear or quadratic surfaces.

linear boundary

Exemplar models
*An exemplar model computes the sum of the similarities between a test exemplar and each of the **stored exemplars** for a given category.*

prototype
*The one **stored exemplar** per category is the mathematical average of the training exemplars of that category.*

roaming-exemplar
*The **stored exemplars** are fitted to best match human behavior.*

RXM[1]
*One **stored exemplar** per category*

RXM[N]
*The same number of **stored exemplars** as training exemplars*

all-exemplar
*All of the training exemplars, regardless of their number, are taken to be **stored exemplars** in the model.*

categorization mechanism

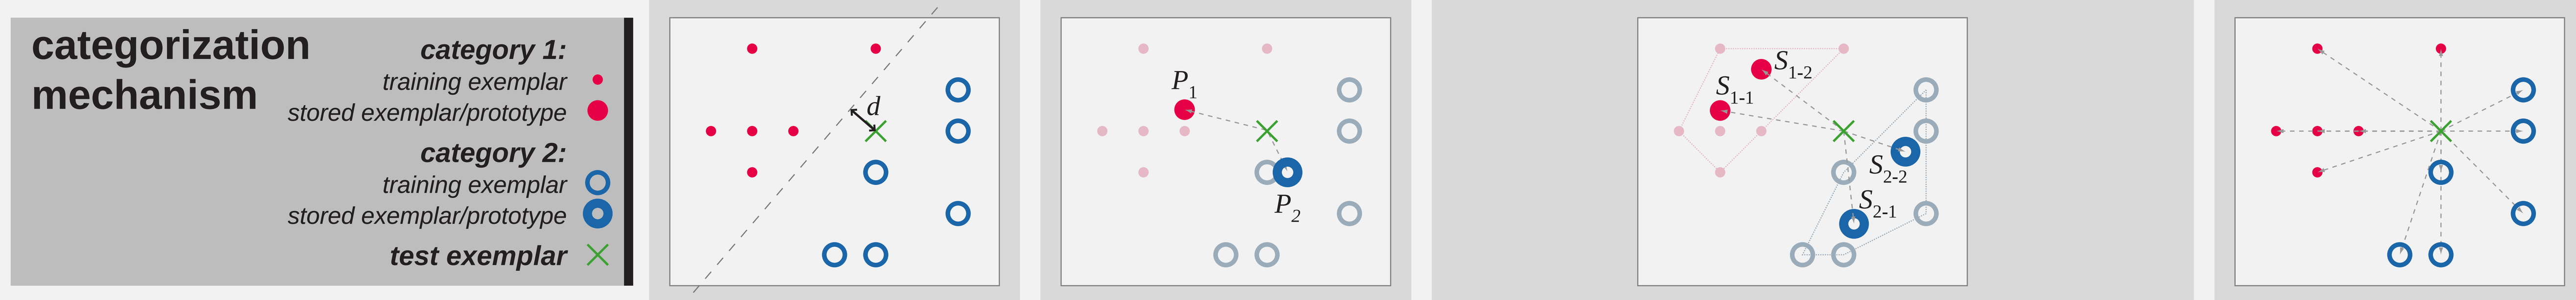

	linear boundary	prototype	RXM[1]	RXM[N]	all-exemplar
stored exemplars	none	1, fixed	1, roaming	*N*, roaming	*N*, fixed
main decision boundary *(orientation)*	arbitrary	constrained	arbitrary	arbitrary	constrained
main decision boundary *(shape)*	plane	plane	plane	curved	curved
iso-probability lines *(shape)*	plane	curved	curved	curved	curved
goodness of fit *rank (AIC)*	**2nd** (177.3)	**4th** (191.0)	**1st** (176.6)	**5th** (279.8)	**3rd** (180.4)

decision surfaces of fitted models
(four examples out of 108 such data sets)

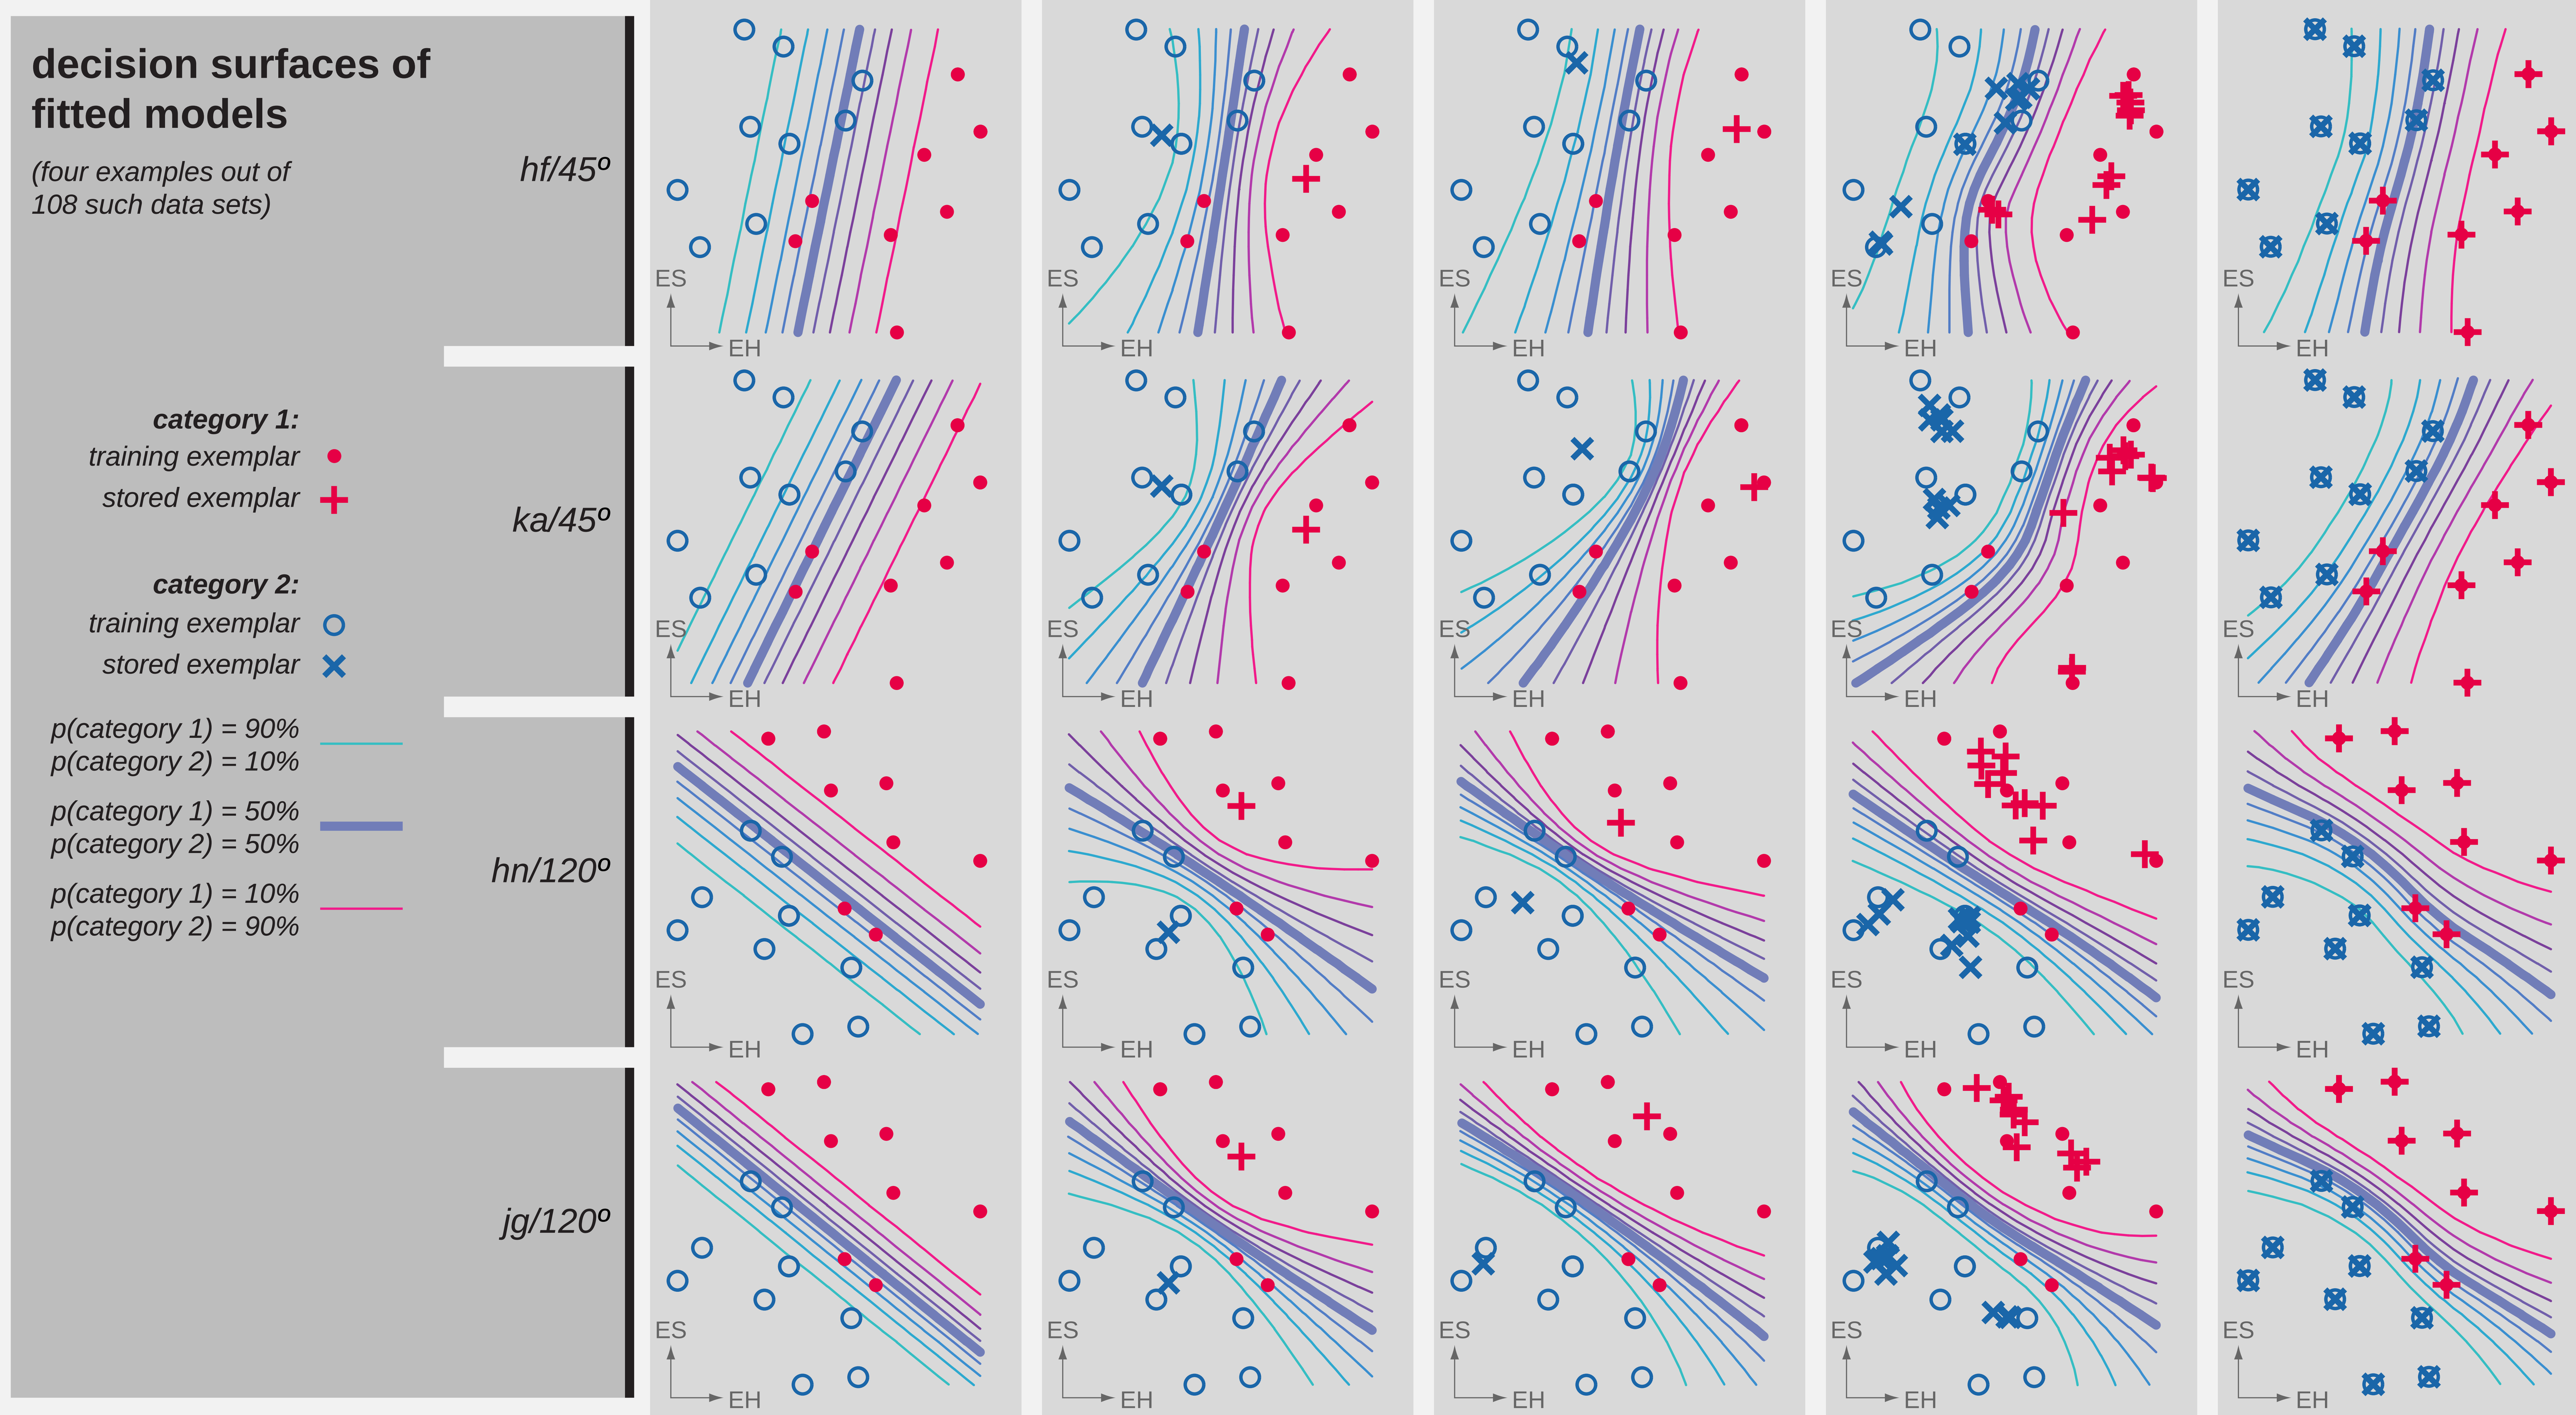

5a Detail: categorization mechanisms

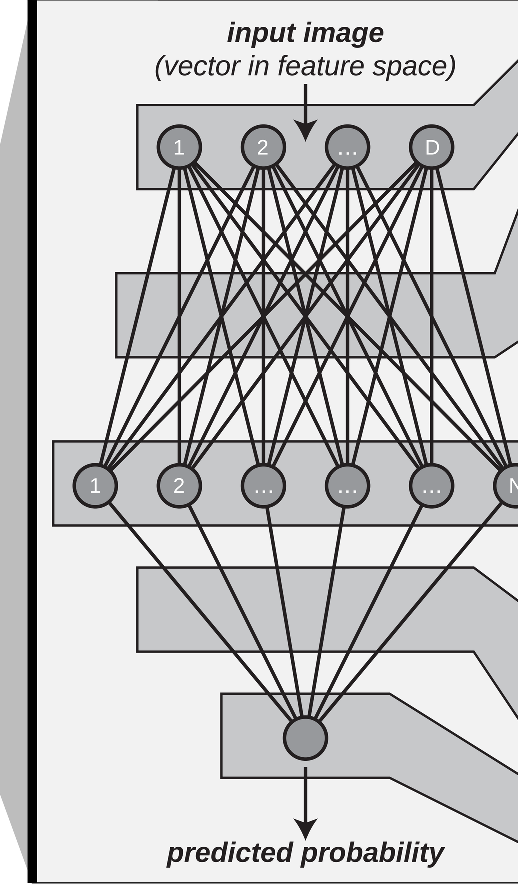

input units $x_1 \ldots x_D$ (eye height, eye separation, ...)

weights $w_{11} \ldots w_{DN}$
In *exemplar models*, these are **attentional weights**, which typically vary only with the input unit, not with the hidden unit.
In *boundary models*, these weights form boundaries in feature space through their dot-products with the input vectors.

hidden units $z_1 \ldots z_N$ (stored exemplars)
In *exemplar models*, the hidden units compute the weighted distance between a stored exemplar $\boldsymbol{\mu}$ and the input pattern $\boldsymbol{x}$:
$$z_j = \mathrm{sigmoid}\left\{\left[\textstyle\sum_D w_{ij}\,(x_i - \mu_i)^2\right]^{1/2}\right\}$$
In *boundary models*, the hidden units compute the distance of input pattern $\boldsymbol{x}$ from the boundary normal to the weight vector $\boldsymbol{w_j}$:
$$z_j = \mathrm{sigmoid}\left(\textstyle\sum_D w_{ij}\,x_i\right)$$

weights $v_0 \ldots v_N$
These weights, through their sign and magnitude, reflect the strength with which each hidden unit is associated with one category or the other. An extra weight v_0 serves as a bias term.

output unit $y = \mathrm{sigmoid}\left(\sum_N v_j z_j\right)$

5b Detail: model fits

Without correcting for free parameters, model fits improves with more memory.

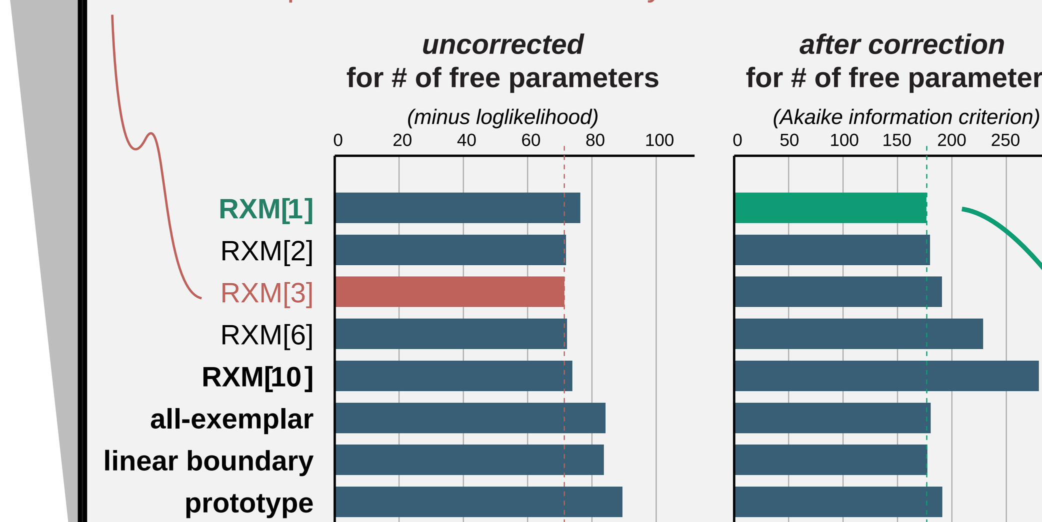

These differences are statistically significant, though small in magnitude.

After correcting for free parameters, ***the RXM[1] fits better than any other model.***

Acknowledgements

This work was supported by a Predoctoral Fellowship from the **Howard Hughes Medical Institute** to R.J. Peters, by the Engineering Research Centers Program of the **National Science Foundation** under Award Number EEC-9402726, by the **National Institute of Mental Health** and by the **W.M. Keck Foundation** Fund for Discovery in Basic Medical Research at Caltech.

Summary

When memory capacity is accounted for with free parameters:

- *a model with **low memory capacity** accounts best for human performance in a subordinate-level categorization task*
- *a successful model with low memory capacity must have **sufficient flexibility in its decision boundary***
- *the success of all-exemplar models (such as the GCM) is due to their relatively more flexible decision boundary, **not** to their high memory capacity*
- *categorization may rely on a **sparse representation** that is different from prototype abstraction*

www.klab.caltech.edu/rjpeters/2001_SFN_Poster.pdf