City Research Online

City, University of London Institutional Repository

An image-driven model for pattern detection, resistant to Birdsall linearisation

Joshua A. Solomon

Centre for Applied Vision Research, City, University of London, EC1V 0HB, United Kingdom

ABSTRACT

If detection were governed by an isolated (and possibly nonlinear) transducer, then a linearisation of the psychometric function (d-prime vs target amplitude) must accompany any threshold elevation due to the addition of external noise. This is the Birdsall theorem. From the fact that noise can elevate threshold *without* linearising the psychometric function, we can safely infer that detection is not governed by an isolated transducer. Heretofore, image-driven models, which accept images or numerical descriptions thereof as input, have proven incompatible with this failure of Birdsall linearisation, unless they incorporate the principle of intrinsic uncertainty, which asserts that detection is governed by the maximum activity in several independent (noisy) sensors. One image-driven model incompatible with the failure of Birdsall linearisation is Watson and Solomon's (*J. Opt. Soc. Am. A*, 14 (1997), 2379) model of visual contrast gain control and pattern masking. Here I report a simple modification – pooling sensor outputs before, instead of after the comparison of input images – allowing that model to predict curved psychometric functions, even when external noise elevates threshold by more than 20 dB, without any detrimental effect to the quality of its fit to pattern-masking thresholds in the absence of noise. The failure of Birdsall linearisation, therefore, does not necessarily imply independent samples of performance-limiting noise in multiple visual sensors. Instead, performance-limiting noise may arise after the visual system combines output from mutually inhibitory sensors.

1. Introduction

The comparison between candidate explanations for visual sensitivity has been greatly facilitated by examining the effects of random perturbations in the visual stimuli (Pelli & Farell, 1999). In the absence of any such external noise, psychometric functions of d′ vs target contrast curve upwards (Nachmias & Sansbury, 1974; Stromeyer & Klein, 1974). An isolated transduction mechanism could produce curvy functions like this, but once external noise had sufficient contrast to impair performance, the psychometric function would have to straighten out. The necessity of this linearisation has been proven mathematically (Lasley & Cohn, 1981). It is known as The Birdsall Theorem.[1]

When replotted as probability correct vs log contrast, linear functions of d′ vs target contrast are well-fit by the relatively shallow Weibull distribution having a shape parameter $\beta = 1.3$, compressed to span the range $(1/m, \ 1)$, where m is the number of alternatives in a multiple-alternative, forced-choice (mAFC) task (May & Solomon, 2013). Despite this specificity, Baker and Meese's (2012) *meta*-analysis of

Birdsall linearisation revealed a wide range of psychometric slopes, when detection was limited by noise whose spatiotemporal frequency spectra were broader than that of the target. It is noteworthy that some of these noises with broad spatiotemporal frequency spectra were nonetheless quite circumscribed in space and time. Indeed, in some cases, the noise's spatial and/or temporal window was identical to that of the target. In such cases the presence of noise can reduce the observer's "intrinsic uncertainty" regarding when and where the target might appear. Although uncertainty reduction can linearise psychometric functions of d′ vs target contrast (Pelli, 1985), it can be avoided by using full-field, dynamic noise, that is "on all the time" (Klein & Levi, 2009).

Solomon and Tyler (2017) measured 45 psychometric functions for detection in relatively long samples of wide-field dynamic noise and another 45 in the absence of noise. Simply put, we found no evidence for Birdsall linearisation. We reported Weibull fits to those 90 psychometric functions. In no noise, the median value of β was 4.0. In high noise the median value was *also* 4.0.[2] Fig. 1 shows a new analysis that tells the same story. To compute these graphs, I re-fit each set of 45 psychometric

---

[1] If two observation intervals generate visual signals i_1 and i_2 before the transducer and o_1 and o_2 after it, then the observer who chooses the larger of the two signals will always choose the same interval either before or after the transducer, since the transducer, being monotonic, cannot alter the signal ordering. Hence the decision on that trial, and every other trial, is unaltered by the transducer, and so performance overall is also unaltered. If enough external noise is added before any transducer nonlinearities, then performance must be linearised since those nonlinearities cannot alter the signal ordering.

[2] With just 88 trials each, it is extremely unlikely that the psychometric functions reported by Solomon and Tyler were artificially flattened by perceptual learning.

functions, assuming various fixed values of β. Fixing β necessarily reduces the likelihood of the fit, but for each data set, that reduction in likelihood was smallest when $\beta = 4$.

If there is no Birdsall linearisation with full-field noise that is on all the time, then no isolated transduction mechanism should be considered an adequate model for detection in noise. On the other hand, if detection were governed by a network of multiple, interacting transduction mechanisms, then The Birdsall Theorem no longer applies. It is conceivable that a such a network could produce realistically curved psychometric functions for detection in noise. The remainder of this paper describes a test of this possibility.

2. Model

Watson and Solomon (1997) described a model of visual contrast gain control and pattern masking that was designed to predict the visibility of compression artefacts in any arbitrary image. I resuscitated that model to see whether it could predict psychometric functions immune to Birdsall linearisation.

The first stage of the model was subdivided into a global contrast sensitivity function (a 2-D log-parabolic filter of spatial frequency), followed by an array of $21,504$ sensors having Gabor-pattern receptive fields.[3] The second stage was contrast gain control. At this stage, sensor outputs were subjected to a power-function transformation. Then the outputs were pooled and used to inhibit (or "normalise") each other. Inhibition strength was defined by Gaussian functions of orientation, spatial frequency, and spatial location. In the third stage of the model, the normalised sensor outputs for one image were subtracted from the corresponding outputs for the other image. These differences were then subjected to Minkowski Pooling.

When the two input images were identical, the Minkowski sum would be zero. Whenever there was a large difference between the two images, the Minkowski sum would be large. "Threshold" image differences could be defined as those for which the Minkowski sum would reliably exceed a random sample from the standard normal distribution. In other words, performance-limiting noise was effectively added at the final, "decision" stage of the model.

To calibrate our model, we used data from Foley and Boynton (1994). In each of the 8 conditions illustrated in Fig. 2, their observers performed a 2-alternative, forced-choice task, in which they had to detect the addition of a low-contrast Gabor pattern to an image containing a Gabor and/or cosine mask. Within each condition Foley and Boynton reported the 92 %-correct detection thresholds for 11 levels of masking contrast (including zero).

3. Model behaviour

Foley and Boynton (1994) reported thresholds for two observers, KMF and JYS. They did not report psychometric slopes, and Watson and Solomon (1997) made no attempt to produce reasonable values for slope when calibrating their model. Nonetheless, computation of those values (i.e., the model's predictions for slope) is straightforward, and those for unmasked detection are far too shallow. Specifically, using the optimal (minimum root-mean-squared error) parameter values for observer KMF's 92 %-correct thresholds, the Weibull distribution best-fitting the psychometric function for unmasked detection has a shape parameter of $\beta = 2.4$. For observer JYS, it has a shape parameter of $\beta = 2.0$. On the other hand, as discussed in the Introduction, psychometric functions for detection are typically well-fit when β is closer to 4.

3.1. New fits 1

Given this restriction ($\beta = 4$) on the psychometric functions for detection, I simultaneously re-fit Watson and Solomon's (1997) model to the 92 %-correct thresholds Foley and Boynton (1994) did report, as well as the (six, unmasked) 58 %-correct thresholds they did not report, which should have been 5.1 dB lower (i.e., 0.255 base-10 log units). Optimal parameter values (*New fits 1* in Table 1) were obtained using Mathematica's FINDMINIMUM routine, with a desired accuracy of 2 significant digits. Addition of these estimated 58 %-correct thresholds effectively penalises shallow psychometric functions for unmasked detection. Consequently, the predicted psychometric slopes increased to
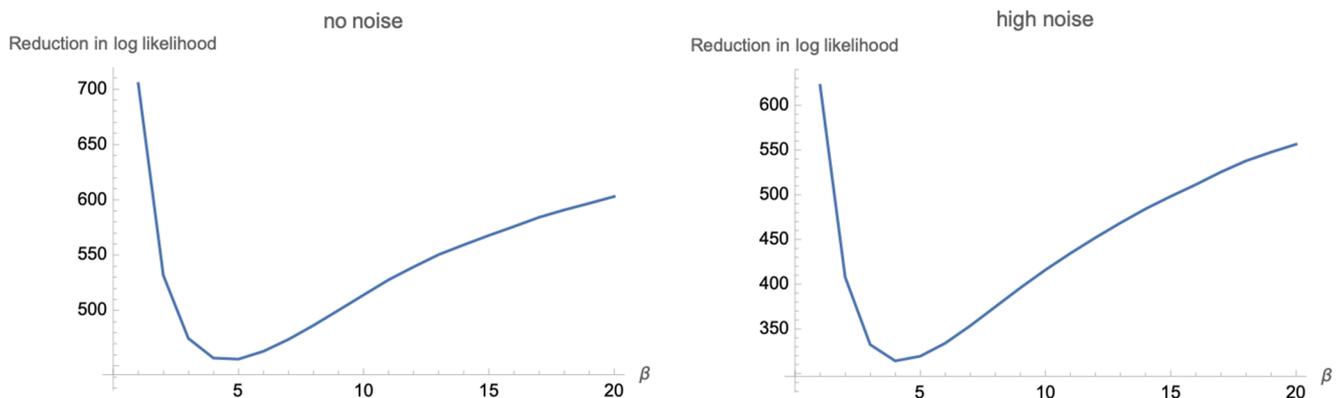

Fig. 1. Reductions in goodness-of-fit to the detection data from Solomon & Tyler (2017). In each condition (no noise and high noise), the Weibull distribution was separately fit to 45 (88-trial) psychometric functions (5 functions/observer × 9 observers) of probability correct vs log contrast. Maximum, joint (base-10) log likelihoods were –295 (no noise) and –199 (high noise). Fixing the Weibull shape parameter β caused likelihoods to fall by the amounts indicated. In both conditions the reduction in likelihood was minimal when β was fixed at the value of 4.

[3] Images were 32 × 32 pixels. 16 Gabor-pattern receptive fields (8 orientations × 2 quadrature phases) were centred on each pixel. The carrier grating in these Gabor patterns had twice the target's spatial frequency. There was an additional array of 16 × 16 × 16 Gabor patterns with carrier gratings matching the target's spatial frequency. Finally (at "the top of the pyramid"), there was an array of 8 × 8 × 16 Gabor patterns with carrier gratings one-half the target's spatial frequency.

$\beta = 2.7$ and $\beta = 2.5$, when using the New-fits-1 parameter values for observers KMF and JYS, respectively.

Two problems arise when attempting to derive the model's predictions for detection in dynamic noise. The first problem is that this model has no temporal components. Consequently, all the simulations reported here were conducted with static stimuli. The second problem arises when the two input images contain different samples of noise. In that case, all the sensor outputs would be different, and the Minkowski
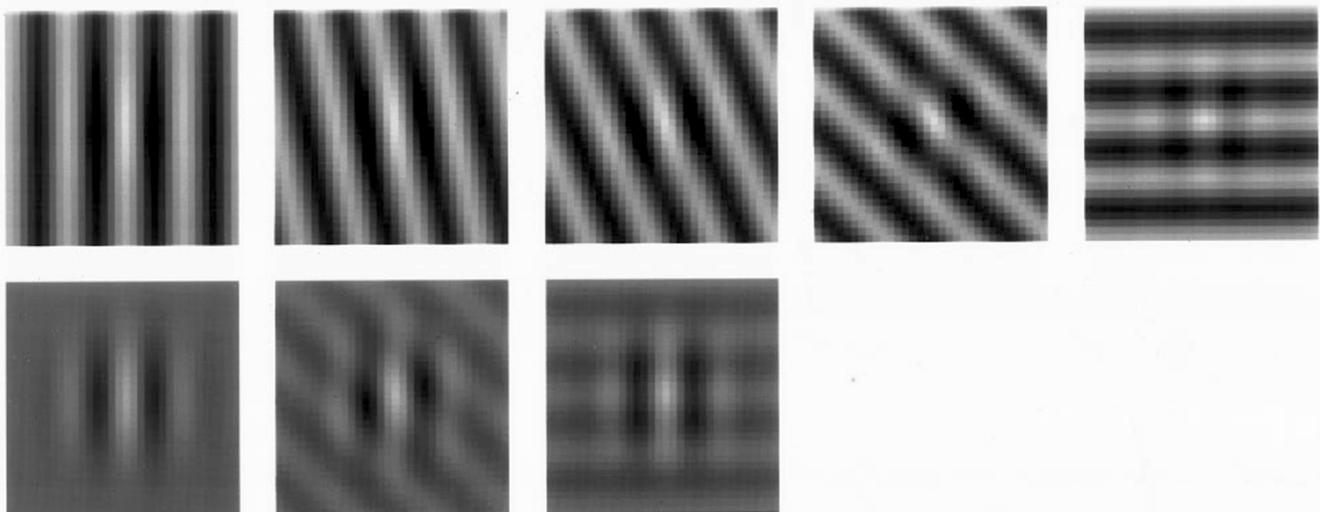
Fig. 2. Example stimuli from Foley and Boynton (1994). The first row shows a Gabor target added to cosine masks at orientations of 0, 11.25, 22.5, 45, and 90 deg. Detection thresholds were reported with each of these masks at 10 logarithmically spaced levels of contrast (plus zero). The second row shows the same target added to an identical Gabor mask, whose contrast was similarly manipulated. The last two panels contain an additional, fixed (at 10 %) contrast cosine mask, at 45 and 90 deg.

Table 1
Estimated Model Parameters and RMS Error for the Two Observers in Foley and Boynton (1994).[1]

Parameter	New fits 1: Six 58 %-correct thresholds for detection		New fits 2: 48 psychometric slopes for detection in noise		New fits 3: optimal sensor only		New fits 4: pooling before comparison	
	KMF	JYS	KMF	JYS	KMF	JYS	KMF	JYS
CSF peak amplitude	39.58	40.63	42.74	42.76	116.3	116.3	49.33	44.00
CSF peak frequency	3.292	4.923	0.982	0.982	0.313	0.324	1.064	1.003
CSF $\log_{10}$ bandwidth*	1.12	1.12	1.12	1.12	1.12	1.12	1.12	1.12
Excitatory exponent	2.312	2.292	2.256	2.272	2.423	2.425	2.314	2.304
Inhibitory exponent*	2	2	2	2	2	2	2	2
Saturation constant	0.0210	0.0171	0.0785	0.0785	0.363	0.337	0.0408	0.0518
Pooling width in $x \times y$[2]	1.55	1.56	0.6846	0.5402	0.5230	0.5300	3.487	3.058
Pooling width in orientation	89.84 deg	81.84 deg	95.89 deg	89.55 deg	91.41 deg	86.86 deg	46.53 deg	41.65 deg
Pooling width in frequency*	0	0	0	0	0	0	0	0
Octave bandwidth of Gabor filters	0.9002	1.065	0.8740	0.8750	0.6994	0.713	2.00	2.00
Minkowski exponent	5.552	5.048	5.493	5.486	N/A	N/A	9.356	9.000
RMS error (original 88 thresholds)	1.647[3] dB	2.040 dB	1.886 dB	2.381 dB	2.030 dB	2.806 dB	1.792 dB	1.955 dB
RMS error (all thresholds)	1.689 dB	2.025 dB	2.510 dB	2.882 dB	2.276 dB	2.656 dB	1.759 dB	1.861 dB

[1] Parameters with asterisks were fixed, as in Watson and Solomon's (1997) calibration.

[2] x and y have units equivalent to nearest-neighbour separation of low-frequency sensors, as in Watson and Solomon's Table 2.

[3] This value is 1.26% lower than that (1.668) reported by Watson and Solomon.

sum could be large even in the absence of any target. My first strategy for dealing with this second problem was to force the two images to have identical samples of noise. Even though I did not have any data on how much this "twinned" noise should elevate threshold, I expected that psychometric functions would remain steep.[4]

3.2. New fits 2

To penalise parameter values that produce shallow psychometric functions, the model's predicted 92 %-correct and 58 %-correct

thresholds were calculated for detection in 6 samples of noise at 8 logarithmically spaced contrasts (root-mean-square values ranging from 0.32 % to 7.9 %). The log-ratio between each pair of thresholds was then subtracted from the predicted value of 0.255. Consequently, each root-mean-squared error quoted for *New fits 2* at the bottom of Table 1 was calculated using these 48 error terms in addition to the 88 differences between the model's predicted thresholds and each observer's data, as reported by Foley and Boynton (1994), plus the six additional error terms for unmasked detection described in *New fits 1*. The 6 samples of noise were held constant, thereby eliminating stochastic fluctuation from the parameter-optimisation routine.

Once the optimal parameters were found, the model was tested with a different sample of twinned noise in a series of 4,096-trial QUEST+ (Watson, 2017) experiments, one for each of 10 logarithmically spaced noise contrasts. Maximum-likelihood values for threshold and psychometric slope are shown in Fig. 3a and b, respectively. These illustrations were assembled using the optimal parameters for observer KMF. The corresponding illustrations (see Supplementary Analyses) for observer JYS were similar. The best fit to data from either observer was obtained when the noise did not elevate threshold until its contrast exceeded

---

[4] Twinned noise has been used in several detection experiments (Swift & Smith, 1983; Burgess & Colborne, 1988; Ahumada & Beard, 1997; Watson, Borthwick, & Taylor, 1997; Beard & Ahumada, 1999; and Solomon, 2002). Some authors (Swift & Smith; Burgess & Colborne; Watson, et al.; cf. Ahumada & Beard) found no appreciable difference between performance with twinned noise and 'random' noise, independently sampled for each interval. None of these authors reported empirical estimates of psychometric slope with twinned noise.
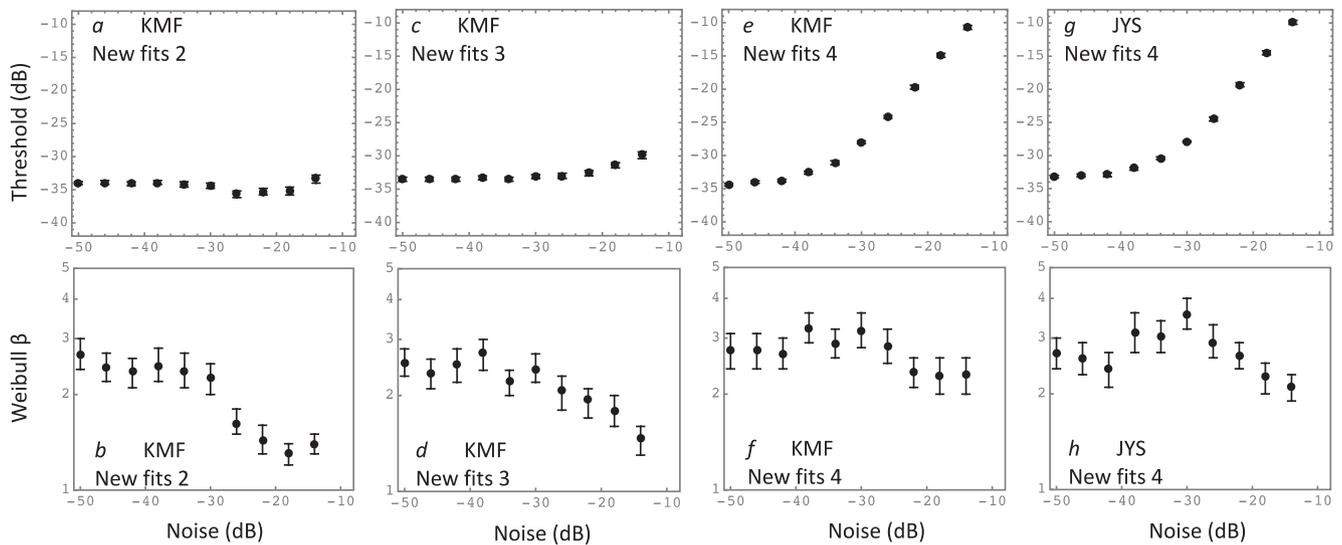
Fig. 3. Simulated detection thresholds and psychometric slopes for a Gabor target in twinned noise. Watson and Solomon's (1997) model (panels *a* and *b*) and various modifications thereof (panels *c – h*) were optimised with data from Foley and Boynton's observers KMF (panels *a – f*) and JYS (panels *g* and *h*), with a penalty for shallow psychometric functions. In all panels, each filled symbol (and error bar) illustrates the maximum-likelihood estimate (and 95% credible interval) derived from a 4,096-trial simulated experiment. In the top row (panels *a*, *c*, *e*, and *g*) most error bars are smaller than symbol size.

about –18 dB, which corresponds to a root-mean-squared contrast of 12 %. However, at even lower contrasts, psychometric slope plummeted to where Weibull β was less than 2.

3.3. New fits 3

Dao, Lu, and Dosher (2006) described a gain-controlled perceptual-template model for pattern detection. Although it shares many features with Watson and Solomon's (1997) model, it has at least one notable difference: no Minkowski summation. Instead, all its decisions are based on the normalised output from a single visual sensor. Inspired by this notable difference, I modified Watson and Solomon's decision rule to ignore output from all sensors except the one best-matched to the target. Optimal parameter values for this modified model were then determined in the same way (i.e., by fitting Foley and Boynton's data with penalties for shallow psychometric functions) they had been determined in *New fits 2*.

This optimal-sensor-only version of Watson and Solomon's (1997) gain-control model proved no more compatible with the failure of Birdsall linearisation than the original version with Minkowski summation. Predicted values of threshold and psychometric slope for observer KMF are shown in Fig. 3c and d, respectively. The corresponding values (see Supplementary Analyses) for observer JYS were similar. Threshold elevation from twinned noise was accompanied by a collapse in psychometric slope.

3.4. New fits 4

An alternative modification of Watson and Solomon's (1997) gain-control model did prove to be somewhat resistant to Birdsall linearisation. This modification was motivated by my desire to make the model compatible with independent samples of noise in the two images. Quite simply, I swopped the subtraction of visual signals with the Minkowski summation; so that the latter came first. More formally, whereas Watson and Solomon's decision statistic can be formulated as.

$$d = \left(\sum | \, _1r_{\bar{\theta}} - \, _2r_{\bar{\theta}} |^M \right)^{1/M} \qquad (1)$$

(cf. their Equation 3), the modified decision statistic can be formulated as.

$$d = \left(\sum | \, _1r_{\bar{\theta}} |^M \right)^{1/M} - \left(\sum | \, _2r_{\bar{\theta}} |^M \right)^{1/M}. \qquad (2)$$

In each of the previous expressions, $_1r_{\bar{\theta}}$ and $_2r_{\bar{\theta}}$ denote (normalised) outputs to the two images from the visual sensor whose preference for spatial frequency, orientation, horizontal position, vertical position, and phase is indexed by the vector $\bar{\theta}$. The Minkowski exponent is *M*. Watson and Solomon's model selected the target image with probability $\Phi(1.4d)$, where $\Phi(x)$ denotes the standard normal (cumulative) distribution of *x*. When the first image contains the target, the modified model also will be correct with probability $\Phi(1.4d)$. When the second image contains the target, it will be correct with probability $\Phi(-1.4d)$.

Optimal parameter values for this modified model were determined in the same way (i.e., with twinned noise) they had been determined in *New fits 2* and *New fits 3*. Thresholds and psychometric slopes for Gabor-pattern detection in twinned noise are shown in Fig. 3e–h. Predictions for detection with independent samples of noise (see Supplementary Analyses) were similar. Note that the modified model simultaneously predicts a plausible threshold elevation from external noise, without a complete collapse in psychometric slope. As the variance of external noise increases from zero, Weibull β increases to a value just above 3 before decreasing to a value just below 3. Thresholds for detection with Gabor and cosine masks are shown in Fig. 4. The fit isn't perfect, but it is pretty satisfying. In particular, the modified model fits the 88 thresholds Foley and Boynton (1994) reported for their observer JYS *better* than the unmodified model, while simultaneously producing plausibly steep psychometric functions for detection in noise.

It isn't entirely clear why the modified decision rule helps the Watson and Solomon (1997) model to resist Birdsall linearisation. In low-to-moderate levels of external noise, decision statistics in both the original (*New fits 2*) and the variant (*New fits 4*) models are dominated by
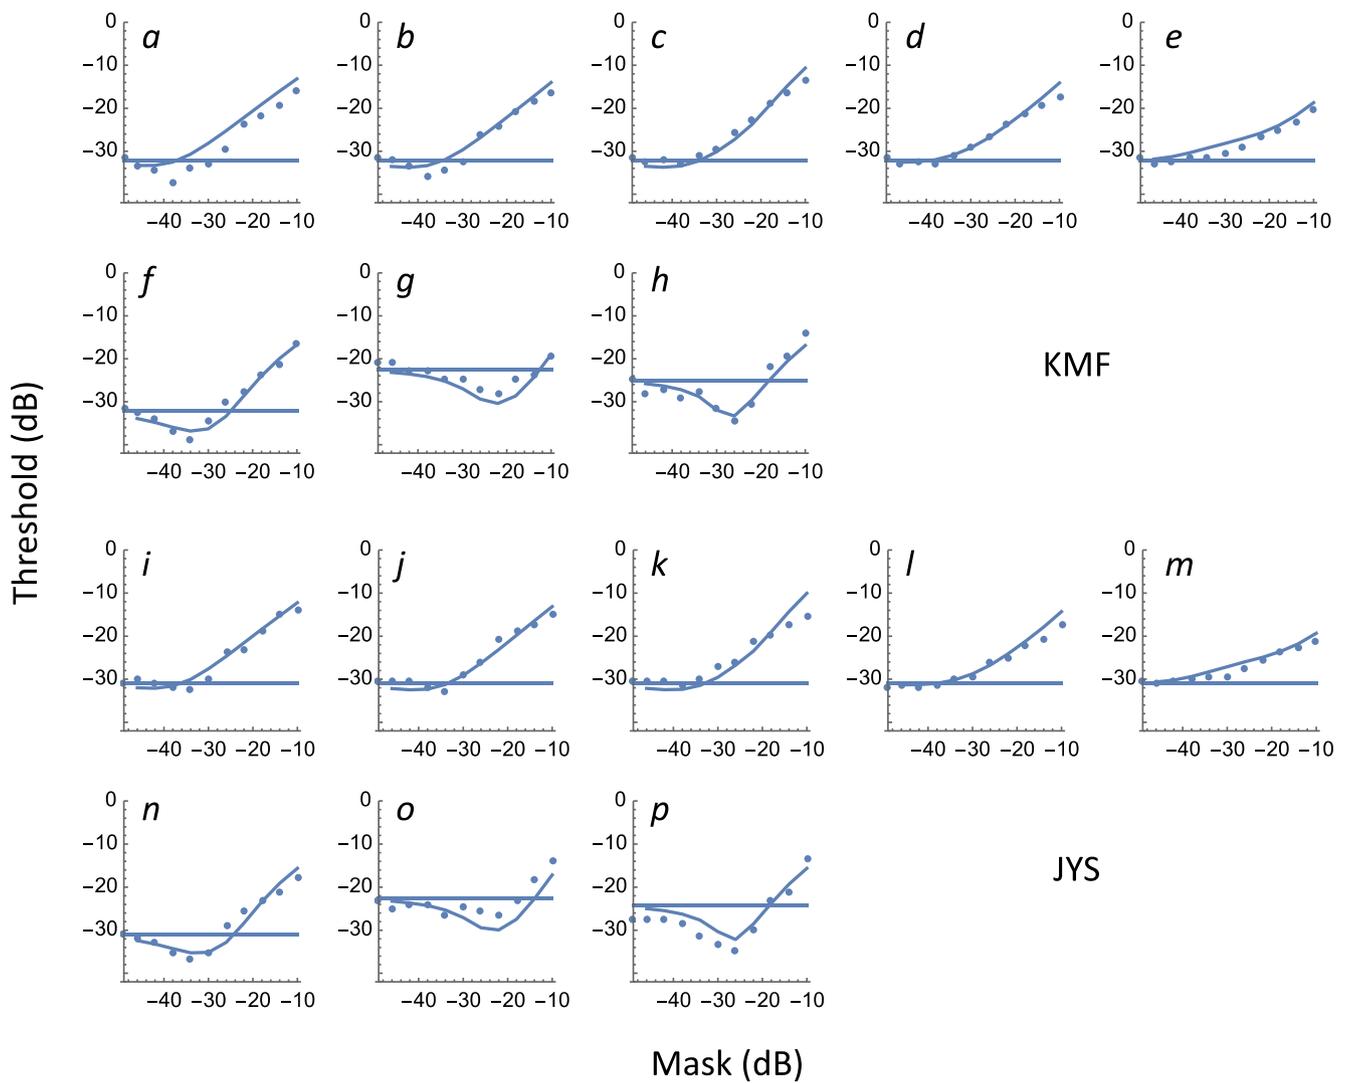
Fig. 4. Thresholds and *New fits 4* for Foley and Boynton's observers KMF (panels *a – h*) and JYS (panels *i – p*). The same parameter values were used for Fig. 3e–h. For each observer, the panel layout corresponds to Fig. 2. Threshold with zero mask contrast (−∞ dB) is plotted on the vertical axis and the model's prediction for this condition is indicated by the horizontal line.

("relevant") sensors tuned to the target.[5] These sensors continue to dominate the original model's decisions when large quantities of twinned noise are added to the Gabor-pattern target. In this case, divisive inhibition is responsible for threshold elevation. In the variant model, on the other hand, many decision statistics in external noise are dominated by arguably irrelevant sensors, tuned to higher frequencies. (See the Supplementary Analyses for illustrations.) Generally speaking, these irrelevant sensors play a greater role in the variant model, making it more similar to intrinsic-uncertainty theory than the original model. In intrinsic-uncertainty theory, psychometric slope is fully determined by the ratio of irrelevant:relevant sensors (Pelli, 1985).

4. Discussion

It does seem sort of embarrassing that, as a field, Psychophysics still hasn't established a generally accepted explanation for the limits of pattern detection. There seems to be two schools of thought regarding why observers make errors in *m*AFC detection tasks. Either they hallucinate targets, and those hallucinations are more intense than the actual targets, or they don't see anything, guess, and guess wrong. Signal-detection theory (Green & Swets, 1966) aligns itself with the first school of thought, but the psychometric functions (probability correct vs log contrast) predicted by simple versions of the theory are too shallow. More complicated alternatives include intrinsic uncertainty, non-linear transduction, and the sensory thresholds favoured by the second school of thought (Solomon, 2007). Of these three alternatives, only intrinsic uncertainty can confer immunity from Birdsall linearisation. The key to this immunity lies in a number of independent detection mechanisms, which I'm calling "sensors" (a.k.a. "micro-analysers"). Closed-form equations can be used to predict *m*AFC performance when sensors are independent (e.g., Solomon & Tyler, 2017). However, many authors have found the requisite number of putatively independent sensors to be implausibly large (>1000), especially if psychometric slope is to remain steep in the presence of external noise. Such Birdsall immunity would require the number of independent sensors to remain

---

[5] Even in the absence of noise and pattern masks, these models behave differently. In the original model, inhibitory signals in or near the sensors best-tuned to the target are greater than those in other sensors. In the variant model, the greatest inhibitory signals are found in sensors tuned to twice the target's frequency. This difference is directly attributable to the variant's relatively large sensor bandwidth. This difference, but not the variant's resistance to Birdsall linearisation, disappears when sensor bandwidth is halved and the other parameters are held constant.

unaffected by the introduction of external noise, and that could happen only if all their receptive fields were mutually orthogonal.

Multiple-sensor models become less implausible when receptive fields are not required to be mutually orthogonal. However, in such cases, no closed-form equations exist that can be used to make quantitative predictions of model performance. Consequently, in this paper, I have examined the behaviour of an image-based model, using Monte Carlo simulations. There would be little reason to struggle with something this complicated if it wasn't potentially immune to Birdsall linearisation.

5. Conclusion

With one small modification to the decision rule, Watson and Solomon's (1997) model of contrast gain control and pattern masking also produces plausible psychometric functions for detection in noise. Of course, this in no way invalidates other models for detection in noise. Solomon and Tyler (2017) championed a model in which detection was always limited by an early source of noise, independently added to the output of each sensor. Indeed, parsimony would suggest early *and* late noises are present whenever sensor outputs are combined. It may prove possible to deduce the relative dominance of these two noises from manipulations of target extent and/or uncertainty.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

I have shared the link to my code in the manuscript itself

Acknowledgments

References

Ahumada, A. J., Jr., & Beard, B. L. (1997). Image discrimination models predict detection in fixed but not random noise. *Journal of the Optical Society of America A, 14,* 2471–2476.

Baker, D. H., & Meese, T. S. (2012). Zero-dimensional noise: The best mask you never saw. *Journal of Vision, 12*(10), 1–12, 20.

Beard, B. L., & Ahumada, A. J., Jr. (1999). Detection in fixed and random noise and parafoveal vision explained by template learning. *Journal of the Optical Society of America A, 16,* 755–763.

Burgess, A. E., & Colborne, B. (1988). Visual signal detection IV. Observer inconsistency. *Journal of the Optical Society of America A, 5,* 617–627.

Dao, D. Y., Lu, Z.-L., & Dosher, B. A. (2006). Adaptation to sine-wave gratings selectively reduces the contrast gain of the adapted stimuli. *Journal of Vision, 6,* 739–759.

Foley, J. M., & Boynton, G. M. (1994). "A new model of human luminance pattern vision mechanisms: Analysis of the effects of pattern orientation, spatial phase and temporal frequency", in Computational Vision Based on Neurobiology, T. B. Lawton, ed. *Proceedings of the SPIE, 2054,* 32–42.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics.* New York: Wiley.

Klein, S. A., & Levi, D. M. (2009). Stochastic model for detection of signals in noise. *Journal of the Optical Society of America A, 26,* B110–B126.

Lasley, D. J., & Cohn, T. E. (1981). Why luminance discrimination may be better than detection. *Vision Research, 21,* 273–278.

Nachmias, J., & Sansbury, R. V. (1974). Grating contrast: Discrimination may be better than detection. *Vision Research, 14,* 1039–1042.

May, K. A., & Solomon, J. A. (2013). Four theorems on the psychometric function. *PLOS ONE, 8,* Article e74815.

Pelli, D. G. (1985). Uncertainty explains many aspects of visual contrast detection and discrimination. *Journal of the Optical Society of America A, 2,* 1508–1532.

Pelli, D. G., & Farell, B. (1999). Why use noise? *Journal of the Optical Society of America A, 16,* 647–653.

Solomon, J. A. (2002). Noise reveals visual mechanisms of detection and discrimination. *Journal of Vision, 2,* 105–120.

Solomon, J. A. (2007). Intrinsic uncertainty explains second responses. *Spatial Vision, 20,* 45–60.

Solomon, J. A., & Tyler, C. W. (2017). Improvement of contrast sensitivity with practice is not compatible with a sensory threshold account. *Journal of the Optical Society of America A, 34,* 870–880.

Stromeyer, C. F., & Klein, S. (1974). Spatial frequency channels in human vision as asymmetric (edge) mechanisms. *Vision Research, 14,* 1409–1420.

Swift, D. J., & Smith, R. A. (1983). Spatial frequency masking and Weber's Law. *Vision Research, 23,* 495–505.

Watson, A. B. (2017). QUEST+: A general multidimensional Bayesian adaptive psychometric method. *Journal of Vision, 17*(3), 1–27, 10.

Watson, A. B., Borthwick, R., & Taylor, M. (1997). Image quality and entropy masking. *SPIE Proceedings, 3016.* paper 1.

Watson, A. B., & Solomon, J. A. (1997). Model of visual contrast gain control and pattern masking. *Journal of the Optical Society of America A, 14,* 2379–2391.