ELSEVIER

Special Issue: Probabilistic models of cognition

Vision as Bayesian inference: analysis by synthesis?

Alan Yuille[1] and Daniel Kersten[2]

[1]Departments of Statistics, UCLA, San Francisco, CA 94115, USA
[2]Department of Psychology, University of Minnesota, Minneapolis, MN 55455-0344, USA

We argue that the study of human vision should be aimed at determining how humans perform natural tasks with natural images. Attempts to understand the phenomenology of vision from artificial stimuli, although worthwhile as a starting point, can lead to faulty generalizations about visual systems, because of the enormous complexity of natural images. Dealing with this complexity is daunting, but Bayesian inference on structured probability distributions offers the ability to design theories of vision that can deal with the complexity of natural images, and that use 'analysis by synthesis' strategies with intriguing similarities to the brain. We examine these strategies using recent examples from computer vision, and outline some important imlications for cognitive science.

Introduction: perception as inference

Experimental studies of biological vision systems have typically been performed using artificial stimuli and tasks. Although much progress has been made, it is questionable how much light these studies shed on the performance of biological vision systems, given the complexity of natural images and natural tasks such as segmentation and object detection and recognition. For example, recent studies suggest that the responses of neurons in V1 to natural stimuli cannot be predicted from their responses to artificial stimuli [1]. It is well-known to computer vision researchers that vision algorithms that work on artificial stimuli almost never generalize to natural images (e.g. there are many algorithms that solve the correspondence problem and estimate depth on Julesz random dot stereograms, but give terrible results when applied to natural images). Using artificial stimuli and tasks to evaluate biological vision is a bit like trying to evaluate the performance of a soldier in battle from his ability at playing with a water pistol.

We argue that the major difficulty of vision arises because natural images are both complex and objectively ambiguous. Similar 3D objects can result in different images, and different objects can result in similar images. Moreover, typical images can be highly complex and consist of hundreds of objects, many of which are overlapping. Therefore theoretical and experimental attempts to understand biological systems must come to terms with the daunting complexity of images somwhat analogous to the way that molecular biologists have had to tame the complexity of the genome.

In this article, we argue that the theoretical principles described in this special issue are rich enough to deal with the complexities and ambiguities of natural images and to perform major perceptual tasks such as recognition and segmentation. The approach is based on Bayesian inference using probability distributions defined on structured representations [2,3]. Vision is treated as an inverse inference problem, in the spirit of Helmholtz [4], where the goal is to estimate the factors that have generated the image. Two major themes follow naturally from this approach.

Firstly, vision as inverse inference presupposes a formalization of how the input is generated, and which of the causes of that input should be estimated. This probabilistic generative approach enables us to define Bayesian Ideal Observers based on the principles of Bayesian Decision Theory and using structured representations and distributions which subsume the standard models of signal detection theory [5] and give theories of optimal performance against which human performance can be compared (the article by Griffiths and Yuille introduces these concepts; see Supplementary material online). Recent survey papers justify this approach and give an entry point to the growing body of literature on this approach [6–8]. An alternative perspective is given by Howe *et al* [9] who question both the use of Bayesian Decision Theory and whether the factors being estimated are physical variables (like reflectance, size; see [4]) or 'orderings of visual stimuli' based on past experience with all such stimuli.

Secondly, this inverse inference perspective suggests that the inference algorithm should combine a top-down generative component with bottom-up processing. The generative component allows the system to internally simulate, or synthesize, from the probability distributions and so is known as 'analysis by synthesis' and relates to the forward and backward projections in the brain [10–16] ([17] for how this may relate to neurotransmitters). But it must be emphasized that analysis by synthesis is not necessarily required for Bayesian inference (work on the

Corresponding authors: Yuille, A. (yuille@stat.ucla.edu) and Kersten, D. (kersten@umn.edu).

Available online 19 June 2006

detection of hands [18]) and so we will give arguments for it in the next section.

The themes in this paper are common to other aspects of cognitive science as described in this special issue (see Conceptual Foundations Editorial by Chater, Tenenbaum and Yuille). The ability to simulate by generative models also occurs naturally within visuo-motor control where they can be used to simulate the consequence of actions (cf. [19–21]). Mental imagery [22] also suggests the ability to do internal simulation, but over long time scales. At a more theoretical level, statistical inference on structured representations offers a common mathematical framework for cognitive science which leads naturally to theoretical models for coupling different sensory modalities and for integrating perception with planning. For example, recent work [23,24] has built on theoretical studies [25,26] to model the coupling of visual cues and haptic cues and shown good fit with experimental data.

The need for bottom-up and top-down processing

We now give arguments why visual inference requires bottom-up and top-down components. This contrasts with standard textbook theories of vision which favour bottom-up processing based on computing low-level representations such as edge maps (cf. [27]) or $2\frac{1}{2}$D representations of depth and shape (cf. [28]).

First observe that low-level vision is ambiguous. For example, it is extremely difficult to determine whether there is an edge present in a small region of an image [29]. These findings are supported by empirical studies of the relative ineffectiveness of edge detectors on natural images [30] and the limitations of regional cues for segmentation [2].

By contrast, high-level vision is rarely ambiguous. The patterns of objects, such as faces or other objects, are complex and rarely occur by chance. Moreover, they are often easy to resolve when they do. For example, patterns in the bark of a tree may occasionally look like a face, but can easily be disambiguated by the alternative explanations for these patterns. This lack of ambiguity for high-level vision is highlighted by the recent successes of computer vision algorithms for detecting high-level objects such as faces and text reliably from natural images [31–33].

From this perspective, a major problem for vision systems is how to use low-level cues to rapidly access the correct high-level models so as to quickly resolve the low-level ambiguities. The difficulty is knowing which high-level model(s) to use. Consider the text-book example of the black and white dalmation dog [34]. Low-level cues for this image contain little evidence to activate a high-level dog model, and so naive subjects take a long time to detect the dog. But subjects who have seen the image before, and know that there is a dog in it, can perceive it instantaneously.

This suggest a visual system where low-level cues make bottom-up proposals which are validated by high-level models. This bottom-up processing should include standard methods such as edge detection and grouping by regional properties – but it must also include special processes for detecting important objects such as faces and text [31,33], and for rapidly classifying the scene [35]. These high-level models access the image, or a filtered version of it, in a top-down process to ensure consistency of the image interpretation. The early visual areas such as V1/V2 are the most natural candidates for such an image representation [11,36]. In certain cases, the bottom-up cues are sufficiently unambiguous and so the object, or scene structure, can be detected without high-level feedback.

Generative models, image parsing and analysis by synthesis

We now describe how this visual system can be implemented. We first give an introductory example which illustrates the main points of the approach on simple examples. Then we describe a more advanced theory that applies to natural images.

Introduction to generative models

First, we consider models for generating an image. Suppose we start with a simple vocabulary of shapes and patterns which contains the letters $A,B,C,\ldots$ We can define a simple probability model for generative images built out of this vocabulary by using templates for each letter and allow the letter to be placed randomly at any position in the image. We can also give an ordering on the letters so that one letter can completely, or partially, occlude another letter. Sampling from this distribution will yield images as shown in Figure 1a.
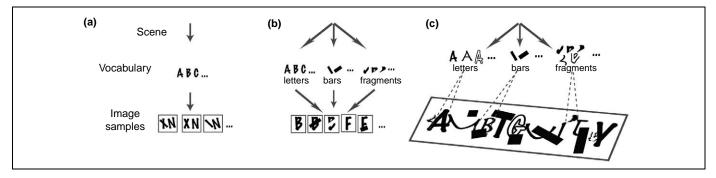

Figure 1. (a) A simple vocabulary for generating the image. There is little or no ambiguity in interpreting images. At worst, the letter X might be confused with a slanted *l* partially occluding a vertical *l*. (b) A richer vocabulary. A given cause, such as a particular letter, can be manifest in many different images. But there are now multiple ways to generate identical images (see text for details). (c) The richer the vocabulary, the greater the image ambiguity, and the harder it is to interpret the image. This leads to a formidable inference problem.

Next, we can make this generative model richer by expanding the vocabulary to include additional objects such as rectangular bars and fragments of letters, Figure 1b. The addition of the elements allows us to generate more complex images but the richness of the vocabulary makes some images potentially ambiguous. The same image can be generated in two distinct ways, see Figure 1b.

This ambiguity in generation leads to potential ambiguities of the inverse process of inference/interpretation. This is resolved by having probabilities on the different ways that images can be generated. The chance alignment of fragments is more likely than a *B* with an invisible (white) occluder (Figure 1b). But a *B* with a visible (black) occluding bar is more likely than a set of fragments which are accidentally aligned with the bar.

We now extend the vocabulary of the generative model in several ways, see Figure 1c. Firstly, we can allow the letters to have properties such as size, font, and shading pattern. Secondly, we can put probabilities on the spatial relations between letters so that, for example, they line up to form words. This extension of the vocabulary leads to further ambiguities.

Analysis by synthesis

But now we turn to a different issue – how fast can we perform inference to determine the most likely way the image was generated? Our explanations, so far, have assumed that we can check all possible ways to generate the image and decide on the most probable. But this becomes completely impractical when the size of the vocabulary becomes very large. So how can we hope to solve the inverse inference problem of estimating which configuration of objects is most likely to have generated the observed image?

We propose an 'analysis by synthesis' strategy where low-level cues, combined with spatial grouping rules (similar to Gestalt laws), make bottom-up proposals which activate hypotheses about objects and scene structures. These hypotheses are accepted, or rejected, by direct comparison with the image (or a filtered version of it) in a top-down process. These bottom-up proposals come with probabilities, which are a measure of their strength. If bottom-up proposals are sufficiently strong (i.e. the low-level cues are

sufficiently unambiguous) then they may be accepted without any need for verification at a lower level.

For example, consider an image containing letters, rectangles, and letter fragments, Figure 2. We can obtain bottom-up cues in several ways. Firstly, we can run an edge detector to discard the shading information (which is often variable) and then spatially group the edges into segments by using the principles of continuity, parallelism and colinearity. This will enable us to get cues for the positions of rectangles and the identity of some letters. These features are used to make bottom-up proposals regarding possible objects. Remaining ambiguity is eliminated in a synthesis stage that tests how well the object models explain the image features. Other more sophisticated cues can be used for the proposals. For example, we can treat text as a type of texture and design detectors which respond to characteristic patterns of text [33]. We stress that this is an illustration of this model only, ignoring shading information may be a mistake for other more realistic stimuli.

Natural image parsing

The acid test for generative models and analysis by synthesis algorithms is whether they can be extended to deal with the complexities of natural images. We now describe recent work which suggests that they can.

This work uses a generative model for images [2,3] which is similar to a two-level probabilistic context free grammar (PCFG) (see article by Griffiths and Yuille; Supplementary material online) and hence we refer to the approach as image parsing. This model is illustrated in Figure 3a where the root node represents the entire image scene. The first level corresponds to the non-terminal nodes representing parameterized models of objects, such as faces or letters, or of generic regions such as textures or shaded patterns. A non-terminal node i has attribute variables (ζ_i, L_i, Θ_i) where ζ_i labels the type of the model (e.g. face, letter, texture or shading), Θ_i denotes the model parameters (e.g. the parameters that determine the shape of the letter), and $R(L_i)$ denotes the region in the image that the model generates. These regions are non-overlapping and the discontinuities across the boundaries are not explicity modeled. Formally, we can summarize the non-terminal nodes by $W = \{(\zeta_i, L_i, \Theta_i): i = 1, \ldots, N\}$ where the number N of nodes is a random variable. These non-
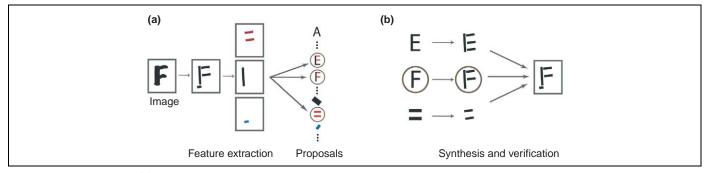

Figure 2. Analysis by synthesis. **(a)** Low-level processing can extract edge features, such as bars, and use conjunctions of these features to make bottom-up proposals to access the higher-level models of objects. **(b)** The high-level objects access the image top-down to validate or reject the bottom-up proposals. In this example, the low-level cues propose that the image can be interpreted as an *E*, an *F*, or a set of parallel bars. But interpreting it as an *F* explains almost all the features in the image and is preferred.
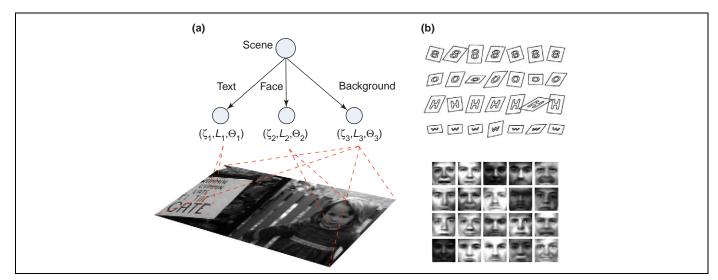
Figure 3. (a) The image is generated by a probabilistic context-free grammar, shown by a two-layer graph with nodes that have properties (ζ, l, θ) corresponding to regions (ζ, l, θ) in the image. (b) Samples from the face model and the letter model, that is, from $p(I_{R(L)} | \zeta, L, \theta)$.

terminal nodes are obtained by sampling from a distribution $P(W) = p(N) \Pi_{i=1}^{N} p(L_i) p(\zeta_i | L_i) p(\Theta_i | \zeta_i)$ (this is conceptually similar to applying production rules to the root node). In turn, the observed intensity values on the image lattice (the terminal nodes of the graph) are obtained by sampling from generative models $p(I_{R(L)} | \zeta, L, \Theta)$ for the specific regions which depend on their model type and their parameters (similar to applying production rules to the non-terminal nodes in a PCFG). This includes models for generating the appearance of faces and letters, see samples in Figure 3b. Overall, this second level gives a model $p(I|W) = \prod_{i=1}^{N} p(I_{R(L_i)} | \zeta_i, L_i, \Theta_i)$ [3]. This combines with the first stage to have a full generative model $p(I|W)p(W)$ for the image.

There are many ways to extend this model by augmenting the number of pattern types, by including Gestalt laws and other principles of spatial organization [37], and by having hierarchical models [38,39]. In particular, the pattern types can be expanded to include material properties which are not explicit objects.

The advantages of a generative model for the entire image include the ability to 'explain away'. Submodels corresponding to different objects, or processes, compete and cooperate to explain different parts of the image (e.g. the letter B plus bar competes with the interpretation of accidentally aligned fragments in Figure 1b). A face model might hallucinate a face in the trunk of a tree; but a tree model can overrule this and provide the correct
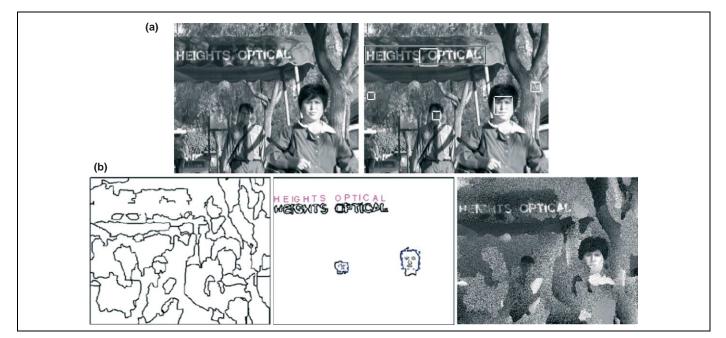

Figure 4. (a) Left: input image. Right: bottom-up proposals for text and faces are shown by boxes. A face is 'hallucinated' in a tree. (b) Left: overall segmentation. Centre: detection of letters and faces. Right: synthesised image.
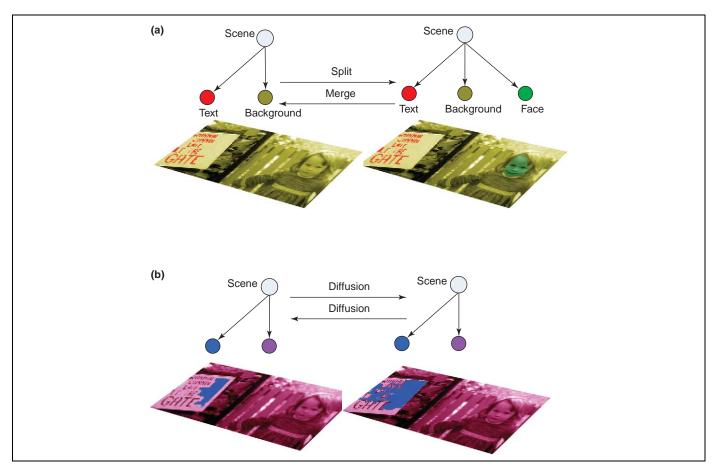
Figure 5. Examples of Markov chain dynamics that change the graph structure or the node attributes of the graph giving rise to different ways to parse the image. Different dynamics, for example creating or deleting nodes, are performed by different kernels *K*.

interpretation of the tree trunk (see Figure 4). In addition, full generative models enforce consistency of the interpretation of the image.

We now switch to the task of performing inference on this generative model to estimate $W^* = \arg\max_W P(W|l)$. This requires a sophisticated inference algorithm that can perform operations such as creating nodes, deleting nodes, diffusing the boundaries, and altering the node attributes.

The strategy used in [3] is to perform analysis by synthesis by a data-driven Markov Chain Monte Carlo (DDMCMC) algorithm. This algorithm is guaranteed to converge by standard properties of MCMC. Informally, low-level cues are used to make hypotheses about the scene which can be verified or rejected by sampling from the models. For example, low-level cues [31,32] can be used to hypothesize that there is a face in a region of the

image. This hypothesis can be validated or rejected by sampling from a generative face model. The bottom-up cues propose that there are faces in the tree bark, but this proposal is rejected by the top-down generative model (Figure 4). Inference is performed by applying a set of operators which change the structure of the parse graph (see Figure 5). These operators are implemented by transition kernels *K* (see Box 1 for a more technical description of the algorithm). The bottom-up cues are based on 'discriminative models' (described in Box 2).

Implications for cognitive science

We claim that the above model for image parsing shares key elements with human visual processing. This claim raises a number of important questions.

Box 2. Generative and discriminative models

Originally discriminative methods were defined by decision rules $\alpha(I)$ which can be described in terms of Bayesian Decision Theory, see box in Griffiths and Yuille's article. These decision rules output discrete values (e.g. 'face' or 'non-face') and there was no attempt to model the probability distribution $P(I,W)$. Discriminative methods of this type include classic techniques like the perceptron and more recent methods such as AdaBoost and Support Vector Machines [49]. More recently, discriminative methods have been generalized to include any method that approximates the posterior distribution $P(W|I)$. Intuitively, these methods make decisions but, by including probabilities, they give a measure of confidence in the decision. This is the sense in which we used discriminative methods in this article. Discriminative methods can be applied to learn approximate distributions $Q(w_1|\phi(I))$ for components w_1 of the full interpretation W, where $\phi(I)$ is a set of features extracted from the image. The key idea, to ensure speed of discriminative proposals, is that the feature $\phi(I)$ can be rapidly extracted and the approximate distribution $Q(W_1|\phi(I))$ is rapid to compute. For example, AdaBoost learning can be used to learn discriminative probabilities for the presence, or absence, of a face at a specific scale, orientation, and location in the image.

Isn't feedback inconsistent with fast processing in human object recognition?

We argue that the bottom-up proposals are consistent with fast feedforward processing. If these proposals are strong, then the high-level percept can occur before top-down validation has begun. There is evidence that reliable diagnostic information for certain categories is available from very simple image measurements [35,32], and that humans make certain categorical decisions sufficiently fast to preclude a verification loop [40,41,42].

Where do the generative models come from?

Ideally the generative models, the discriminative models, and the stochastic grammar would all be learnt from natural images. This is not difficult in principle because, as discussed in Griffiths and Yuille (see Supplementary material online), learning the model from data is simply another example of statistical inference. The Helmholtz machine [43] gives an illustration of how a generative model, and an inference algorithm, can be learnt. This approach, however, has been applied only to simple visual stimuli. Similarly Friston [16] suggests learning models using the Expectation–Maximization algorithm. Although this is a useful metaphor, the challenge is to see whether this idea can be translated to algorithms that can deal with the complexities of natural images.

Learning generative and discriminative models is an extremely difficult problem in practice owing to the large dimensionality of natural images. There has recently, however, been dramatic progress on the similar, but arguably simpler, problem of learning a stochastic grammar for natural languages (see Chater and Manning, this issue [44]). At present, different components of the image parsing model are learnt individually. For example, the discriminative models for text and faces are trained using labelled examples of 'face', 'text', and 'non-face', 'non-text'. Similarly the generative models for faces and text are learnt from labelled examples of faces and text. This, however, is far easier than unsupervised learning of the full stochastic grammar.

Where is attention in the image parsing model?

The current DDMCMC algorithm contains no explicit attention mechanism. Instead bottom-up cues act everywhere in parallel to activate the generative models. But there are some mechanisms in DDMCMC which are similar to attention – for example, proposals which have high confidence are processed quickly. Moreover, it is possible to add an attentional mechanism whereby processing is directed to maximize criteria such as the expected gain in information. In other words, uncertainty in the residual signal helps prioritize where DDMCMC should process the image. This could be part of the mechanisms for driving extrinsic attention. It is unclear how this relates to standard theories of attention and the pre/post attentive distinction [27]. But it is also unclear how these types of theories, developed for artificial idealized stimuli, relate to the performance of human observers on realistic natural images. For example, Li *et al.* [45] show that human subjects can pre-attentively detect complex objects, such as animals, in natural images.

How does this relate to neural mechanisms?

There have long been speculations about the relationship between analysis by synthesis and forward and backward pathways in the cortex [11]. For a recent review of the ascending and descending pathways between visual areas, see Ref. [15].

Although it is impractical at present to test the details of analysis by synthesis models there are some relevent findings using functional magnetic resonance imaging (fMRI) and high-density electrical mapping of evoked response potentials (ERP). Earlier fMRI work by a number of groups has shown that the human lateral occipital complex (LOC) increases activity during the perception of object completion. Murray *et al.* [46] used fMRI to show that when local visual information is perceptually organized into whole objects, activity in human primary visual cortex (V1) decreases over the same approximate period that activity in higher, lateral occipital areas (LOC) increases. The authors interpret the activity changes in terms of high-level hypotheses that compete to explain away the incoming retinal data. This interpretation is also consistent with findings of another group who combined ERP and fMRI [47]. Based on timing measurements, they concluded that dorsal regions provide input to LOC, where signatures of illusory contour completion first appear, then followed by activity in V1/V2. Activity in fusiform gyrus has also been associated with object recognition. Bar *et al.* [48] used magnetoencephalography (MEG) and fMRI in an object recognition task to show that low spatial frequency specific activity appeared in left orbitofrontal cortex 50 ms earlier than in fusiform gyrus of temporal cortex. They interpreted their results as consistent with a fast, spatially coarse analysis

that selects probable object interpretations which are subsequently integrated with bottom-up information.

Conclusion

We have argued that the major goal of vision science should be to determine how biological systems work under natural conditions and when performing natural tasks. This requires understanding and modeling the complexity of images. We argue that studies of perceptual abilities on simple synthetic stimuli can be misleading and unrepresentative of how perceptual systems function under realistic conditions.

By using recent examples from the computer vision literature, we showed that probability distributions defined on structured representations offer the promise to model natural images. Current research is extending this model including more sophisticated representations, better bottom-up cues, and extensions to greater ranges of objects and scenes (see also Editorial 'Where next?' in this issue). One advantage of this type of theory is that is readily 'extensible' in the sense that by designing increasingly sophisticated generative models of this form we can in principle develop artificial visual systems of arbitrary effectiveness. These generative models can be used as Bayesian Ideal Observers.

The inference algorithm for this generative model is intriguing because it naturally follows the 'analysis by synthesis' strategy that may correspond to the forward and backward pathways in the cortex. This algorithm combines bottom-up and top-down processing by using low-level cues to activate high-level models which are compared with the image in a top-down process. It is hoped that as techniques like fMRI, EEG, MEG and multi-unit recording continue to develop it will be possible to make direct experimental predictions from these models.

Supplementary data

Supplementary data associated with this article can be found at doi:10.1016/j.tics.2006.05.002

References

1 Carandini, M. *et al*. (2005) Do we know what the early visual system does? *J. Neurosci.* 25, 10577–10597

2 Tu, Z. and Zhu, S.C. (2002) Image segmentation by data-driven Markov Chain Monte Carlo. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 657–673

3 Tu, Z. *et al*. (2005) Image parsing: Unifying segmentation, detection and recognition Int. *J. Comput. Vis.* 63, 113–140

4 Helmholtz, H. (1867) *Handbuch der Physiologischen Optik*, Voss, Leipzig

5 Green, D.M. and Swets, J.A. (1974) *Signal Detection Theory and Psychophysics*, Krieger Publishing Company, New York

6 Mamassian, P. *et al*. (2002) Bayesian modelling of visual perception. In *Probabilistic Models of the Brain: Perception and Neural Function* (Rao, R.P.N. *et al*., eds), pp. 13–36, MIT Press

7 Geisler, W.S. and Kersten, D. (2002) Illusions, perception and bayes. *Nat. Neurosci.* 5, 508–510

8 Kersten, D. *et al*. (2004) Object perception as Bayesian inference. *Annu. Rev. Psychol.* 55, 271–304

9 Howe, C.Q. *et al*. (2006) Comparison of Bayesian and empirical ranking approaches to visual perception. *J. Theor. Biol.*

10 MacKay, D. (1956) Towards an information-flow model of human behavior. *Br. J. Psychol.* 47, 30–43

11 Mumford, D. (1992) On the computational architecture of the neocortex. ii. the role of cortico-cortical loops. *Biol. Cybern.* 66, 241–251

12 Ullman, S. (1996) *High-Level Vision: Object Recognition and Visual Cognition*, MIT Press

13 Grenander, U. (1996) *Elements of Pattern Theory (Johns Hopkins Series in the Mathematical Sciences)*, Johns Hopkins University Press

14 Rao, R.P. and Ballard, D.H. (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87

15 Friston, K. (2003) Learning and inference in the brain. *Neural Netw.* 16, 1325–1352

16 Friston, K. (2005) A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 360, 815–836

17 Dayan, P. and Yu, A. (2005) Norepinephrine and neural interrupts NIPS advances in neural information processing systems vol. 18

18 Coughlan, J. *et al*. (2000) Efficient deformable template detection and localization without user initialization. *Comput. Vis. Image Underst.* 78, 303–319

19 Kawato, M. (1999) Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* 9, 718–727

20 Todorov, E. and Jordan, M.I. (2002) Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5, 1226–1235

21 Flanagan, J.R. *et al*. (2003) Prediction precedes control in motor learning. *Curr. Biol.* 13, 146–150

22 Klein, I. *et al*. (2000) Transient activity in the human calcarine cortex during visual-mental imagery: an event-related fmri study. *J. Cogn. Neurosci.* 12(Suppl 2), 15–23

23 Ernst, M.O. and Banks, M.S. (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–433

24 Schrater, P.R. and Kersten, D. (2000) How optimal depth cue integration depends on the task. *Int. J. Comput. Vis.* 40, 73–91

25 Clark, J.J. and Yuille, A.L. (1990) *Data Fusion for Sensory Information Processing Systems (Kluwer International Series in Engineering and Computer Science, Vol. 105: Robotics)*, Kluwer Academic Publishers

26 Yuille, A.L. and Bülthoff, H.H. (1996) Bayesian decision theory and psychophysics. In *Perception as Bayesian Inference* (Knill, D. and Richards, W., eds), pp. 123–161, Cambridge University Press

27 Palmer, S.E. (1999) *Vision Science: Photons to Phenomenology*, MIT Press

28 Marr, D. (1982) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*, W.H. Freeman & Co.

29 McDermott, J. (2004) Psychophysics with junctions in real images. *Perception* 33, 1101–1127

30 Konishi, S.M. *et al*. (2003) Statistical edge detection: Learning and evaluating edge cues. *IEEE Trans. Pattern Anal. Mach. Intell.* 25, 57–74

31 Viola, P. and Jones, M.J. (2004) Robust real-time face detection. *Int. J. Comput. Vis.* 57, 137–154

32 Ullman, S. *et al*. (2002) Visual features of intermediate complexity and their use in classification. *Nat. Neurosci.* 5, 682–687

33 Chen, X. and Yuille, A. (2004) Detecting and reading text in natural scenes. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (CVPR'04), pp. 366–373, IEEE

34 Gregory, R.L. (1973) *The Intelligent Eye* (p. 14), McGraw-Hill

35 Oliva, A. and Torralba, A. (2001) Modeling the shape of the scene: A holistic representation of the spatial envelope. *Int. J. Comput. Vis.* 42, 145–175

36 Lee, T.S. and Mumford, D. (2003) Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 20, 1434–1448

37 Tu, Z. and Zhu, S.C. (2002) Parsing images into region and curve processes. In *Lecture Notes in Computer Science: Lecture Notes in Artificial Intelligence* (Vol. 2352), pp. 397–407, Springer-Verlag

38 Ahissar, M. and Hochstein, S. (2004) The reverse hierarchy theory of visual perceptual learning. *Trends Cogn. Sci.* 8, 457–464

39 Feng, H. and Zhu, S. (2005) Bottom-up/top-down image parsing by attribute graph grammar. In *Proc. 10th IEEE Int. Conf. Comput. Vis.* (ICCV'05), pp. 1778–1785, IEEE

40 VanRullen, R. and Thorpe, S.J. (2001) The time course of visual processing: from early perception to decision-making. *J. Cogn. Neurosci.* 13, 454–461

41 Foxe, J.J. and Simpson, G.V. (2002) Flow of activation from V1 to frontal cortex in humans. a framework for defining 'early' visual processing. *Exp. Brain Res.* 142, 139–150

42 Johnson, J.S. and Olshausen, B.A. (2003) Timecourse of neural signatures of object recognition. *J. Vis.* 3, 499–512

43 Dayan, P. *et al*. (1995) The Helmholtz machine. *Neural Comput.* 7, 889–904

44 Chater, N. and Manning, C.D. (2006) Probabilistic models of language processing and acquisition. *Trends Cogn. Sci.* DOI:10.1016/j.tics.2006.05.006

45 Li, F.F. *et al*. (2002) Rapid natural scene categorization in the near absence of attention. *Proc. Natl. Acad. Sci. U. S. A.* 99, 9596–9601

46 Murray, S.O. *et al*. (2002) Shape perception reduces activity in human primary visual cortex. *Proc. Natl. Acad. Sci. U. S. A.* 99, 15164–15169

47 Murray, M.M. *et al*. (2004) Setting boundaries: brain dynamics of modal and amodal illusory shape completion in humans. *J. Neurosci.* 24, 6898–6903

48 Bar, M. *et al*. (2006) Top-down facilitation of visual recognition. *Proc. Natl. Acad. Sci. U. S. A.* 103, 449–454

49 Duda, R.O. *et al*. (2001) *Pattern Classification*, 2nd edn, Wiley