

Formal directed graphical models

- ▶ *Directed graphical models* are formally specified as follows. The random variables X_μ are defined at the nodes $\mu \in \mathcal{V}$ of a graph.
- ▶ The edges \mathcal{E} specify which variables directly influence each other. For any node $\mu \in \mathcal{V}$, the set of parent nodes $pa(\mu)$ are the set of all nodes $\nu \in \mathcal{V}$ such that $(\mu, \nu) \in \mathcal{E}$, where (μ, ν) means that there is an edge between nodes μ and ν pointing to node μ . We denote the state of the parent node by $\vec{X}_{pa(\mu)}$.
- ▶ This gives a local *Markov property* – the conditional distribution $P(X_\mu | \vec{X}_{/\mu}) = P(X_\mu | \vec{X}_{pa(\mu)})$, so the state of X_μ is directly influenced only by the state of its parents (note $\vec{X}_{/\mu}$ denotes the states of all nodes except for node μ). Then the full distribution for all the variables can be expressed as:

$$P(\{X_\mu : \mu \in \mathcal{V}\}) = \prod_{\mu \in \mathcal{V}} P(X_\mu | \vec{X}_{pa(\mu)}). \quad (43)$$

Directed graphical models: Divisive normalization and Bayes-Kalman

- ▶ Two examples of directed graphical models are described in the Yuille-Kersten book chapter.
- ▶ First, when we studied divisive normalization used to represent the dependencies between the stimuli, the filter responses, and the common factor.
- ▶ Second, when exploring the Bayes-Kalman filter, where the hidden state x_t at time t “causes” the hidden state x_{t+1} at time t and the observation y_t . The Bayes-Kalman filter will be discussed later in this lecture.
- ▶ Note that in some situations, the directions of the edges indicate physical causation between variables, but in others, the arrows merely represent statistical dependence. The relationship between graphical models and causality is complex and is clarified in (Pearl, 2000).

Causal structure: Taxonomy of cue interactions

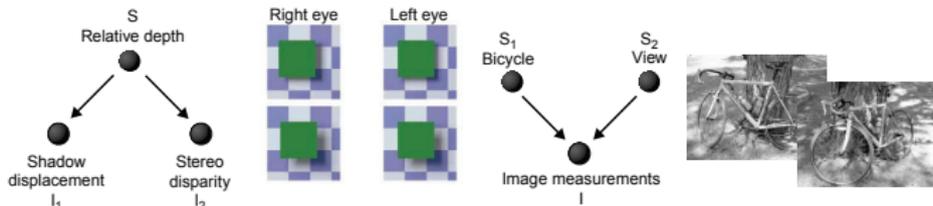


Figure 36: Graphical models give a taxonomy of different ways in which visual cues can be combined. Left: An example of common cause. The shadow and binocular stereo cues are caused by the same event – two surfaces with one partially occluding the other. Right: The image of the bicycle is caused by the pose of the bicycle, the viewpoint of the camera, and the lighting conditions.

Graphical models and explaining away (I)

- ▶ Graphical models can be used (Pearl, 1988) to illustrate the phenomena of *explaining away*. This describes how our interpretations of events can change suddenly as new information becomes available.
- ▶ For example, suppose a house alarm A can be activated by either a burglary B or by an earthquake E . This can be modeled by $P(A|B, E)$ and priors $P(B), P(E)$ for a burglary and an earthquake. In general, the prior probability of a burglary is much higher than the prior probability of an earthquake. So if an alarm goes off, then it is much more probable to be caused by a burglary, formally $P(B|A) \gg P(E|A)$. But suppose, after the alarm has sounded, you are worried about your house and check the Radio only to discover that there has been an earthquake. In this case, this new information “explains away” the alarm, so you stop worrying about a burglary.

Graphical models and explaining away (I)(A)

- ▶ The burglary example can be represented by a graph with nodes A, B, E, R . Each node has a binary-valued state variable X_A, X_B, X_E, X_R . If $X_A = 1$ the alarm activated ($X_A = 0$ if it did not). Similarly the variable X_B, X_E indicate whether there is a burglary, an earthquake or not. $X_R = 1$ if there is radio news about an earthquake ($X_R = 0$ if there is not).
- ▶ The joint distribution $P(X_A, X_B, X_E, X_R)$ can be decomposed as $P(X_A|X_B, X_E)P(X_B)P(X_E)P(X_R|X_E)$. This decomposition isolates the causal factors (but see Pearl 2000, true causality requires the ability to intervene and perform "graph surgery"). It also reduces the amount of data required to model the problem, both for learning the distributions and for inferring the best interpretation X_B, X_E of the observed data X_A, X_R .
- ▶ If there is no radio, then the graph is simplified $P(X_A, X_B, X_E) = P(X_A|X_B, X_E)P(X_B)P(X_E)$. Without the evidence from the radio we will probably interpret the alarm firing as a burglary. The new evidence (from the radio) changes our interpretation. Pearl invented this example to argue for the importance of using probabilities to model reasoning, because non-probabilistic approaches like conventional logic would find it difficult to deal with this situation.

Graphical models and explaining away (I)(B)

- ▶ Learning and Computation. How many parameters does the simple model $P(X_A, X_B, X_E)$ need? How do we compute the posteriors $P(X_B, X_E|X_A)$?
- ▶ A general distribution $P(X_A, X_B, X_E)$ has $2^3 - 1 = 7$ parameters (each variable X takes two possible values, yielding 2^3 but we subtract 1 because of the constraint $\sum_{X_A, X_B, X_E} P(X_A, X_B, X_E) = 1$). But if $P(X_A, X_B, X_E) = P(X_A|X_B, X_E)P(X_B)P(X_E)$ then it takes only $4 + 1 = 5 = 6$ parameters (note $P(X_A|X_B, X_E)$ has 4 parameters, because $\sum_{X_A} P(X_A|X_B, X_E) = 1$ for each possible value of X_B, X_E , and there are four possible values for X_B, X_E . This means that a (little) less data is required to learn it.
- ▶ To estimate the probability of an earthquake or a burglary we must compute the posterior distribution $P(X_B, X_E|X_A)$. This is $P(X_B, X_E|X_A) = \frac{P(X_A, X_B, X_E)}{P(X_A)} = \frac{P(X_A, X_B, X_E)}{\sum_{X_B, X_E} P(X_A, X_B, X_E)} = \frac{P(X_A|X_B, X_E)P(X_B)P(X_E)}{\sum_{X_B, X_E} P(X_A|X_B, X_E)P(X_B)P(X_E)}$.
- ▶ To compute the posteriors for earthquakes and burglaries separately, we compute $P(X_E|X_A) = \sum_{X_B} P(X_B, X_E|X_A)$ and $P(X_B|X_A) = \sum_{X_E} P(X_B, X_E|X_A)$. (The general rule is "sum out variables you are not interested in).

Graphical models and explaining away (I)(C)

- ▶ Similarly for the full model

$P(X_A, X_B, X_E, X_R) = P(X_A|X_B, X_E)P(X_R|X_E)P(X_B)P(X_E)$, we find that there are $4 + 2 + 1 + 1 = 8$ parameters (instead of $2^4 - 1 = 15$).

- ▶ The posteriors $P(X_B, X_E|X_A, X_R)$ are computed to be

$$\frac{P(X_A|X_B, X_E)P(X_R|X_E)P(X_B)P(X_E)}{\sum_{X_B, X_E} P(X_A|X_B, X_E)P(X_R|X_E)P(X_B)P(X_E)}.$$

Graphical models and explaining away (II)

- ▶ Variants of this phenomena arise in vision. Suppose you see the “partly occluded T ” where a large part of the letter T is missing. In this case there is no obvious reason that part of the T is missing, so the perception may be only of two isolated segments. On the other hand, if there is a grey smudge over the missing part of the T , then most observers perceive the T directly. The presence of the smudge “explains away” why part of the T is missing.
- ▶ The Kanizsa triangle can also be thought of in these terms. The perception is of three circles partly occluded by the triangle. Hence the triangle explains why the circles are not complete. We will give a closely related explanation when we discuss model selection.

Directed graphical models and visual tasks (I)

- ▶ The human visual system performs a range of visual tasks, and the way cues are combined can depend on the tasks being performed.
- ▶ For example, consider determining the shape of a shaded surface. In most cases we need only shape from shading to estimate the shape of the surface. But occasionally we may want to estimate the light source direction.
- ▶ This can be formulated by a model $P(I|S, L)P(S), P(L)$, where I is the observed image, S is the surface shape, and L is the light source direction. $P(I|S, L)$ is the probability of generating an image I from shape S with lighting L , and $P(S), P(L)$ are prior probabilities on the surface shape and the lighting.

Directed graphical models and visual tasks (II)

- ▶ If we only want to estimate the surface shape S , then we do not care about the lighting L . The optimal Bayesian procedure is to integrate it out to obtain a likelihood $P(I|S) = \int dL P(I|S, L) P(L)$, which is combined with a prior $P(S)$ to estimate S .
- ▶ Conversely, if we only want to estimate the lighting, then we should integrate out the surface shape to obtain a likelihood $P(I|L) = \int dS P(I|S, L) P(S)$ and combine it with a prior $P(L)$.
- ▶ If we want to estimate both the surface shape and the lighting, then we should estimate them using the full model $P(I|S, L)$ with priors $P(S)$ and $P(L)$.
- ▶ “Integrating out” nuisance, or generic, variables relates to the *generic viewpoint assumption* (Freeman, 1994) which states that the estimation of one variable, such as the surface shape, should be insensitive to small changes in another variable (e.g., the lighting).

Model selection.

- ▶ Certain types of cue coupling require *model selection*.
- ▶ While some cues, such as binocular stereo and motion, are usually valid in most places of the image, other cues are only valid for subparts of each image. For example, the lighting and geometry in most images are too complex to make shape from shading a reliable cue. Also shape from texture is only valid in restricted situations.
- ▶ Similarly, the visual system can use *perspective cues* to exploit the regular geometrical structure in the ball-in-a-box experiments. But such cues are only present in restricted classes of scenes, which obey the “Manhattan world” assumption. These cues will not work in the jungle. These considerations show that cue combination often requires *model selection* in order to determine in what parts of the image, if any, the cues are valid.

Model selection illustration

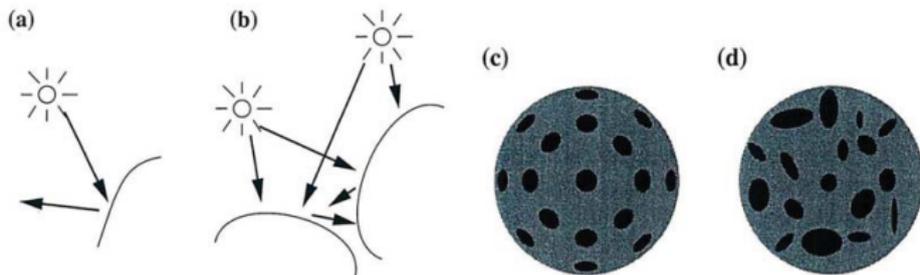


Figure 37: Model selection may need to be applied to decide if a cue can be used. Shape from shading cues will work for case (a) because the shading pattern is simply due to a smooth convex surface illuminated by a single source. But for case (b) the shading pattern is complex – due to mutual reflection between the two surfaces – and so shape from shading cues will be almost impossible to use. Similarly, shape from texture is possible for case (c), because the surface contains a regular texture pattern, but is much harder for case (d), because the texture is irregular.

Model selection examples

- ▶ Model selection also arises when there are several alternative ways to generate the image.
- ▶ By careful experimental design, it is possible to adjust the image so that small changes shift the balance between one interpretation and another.
- ▶ Examples include the experiments with two rotating planes that can be arranged to have two competing explanations (Kersten et al., 1992). With slight variations to the transparency cues, the two surfaces can be seen to move rigidly together or to move independently (see <http://youtu.be/gSrUBpovQdU>).

Model selection: shadows and specularity

- ▶ A classic experiment (Blake & Bulthoff, 1990) studies human perception using a sphere with a Lambertian (diffuse) reflection function, which is viewed binocularly.
- ▶ A specular component is adjusted so that it can lie in front of the sphere, between the center and the sphere, or at the center of the sphere.
- ▶ If the specularity lies at the center, then it is perceived to be a transparent light bulb.
- ▶ If the specularity is placed between the center and the sphere, then the sphere is perceived to be shiny and specular.
- ▶ If the specularity lies in front of the sphere, then it is perceived as a cloud floating in front of a matte (Lambertian sphere).
- ▶ This is interpreted as strong coupling using model selection (Yuille & Bulthoff, 1996).

Model selection examples: Illustration

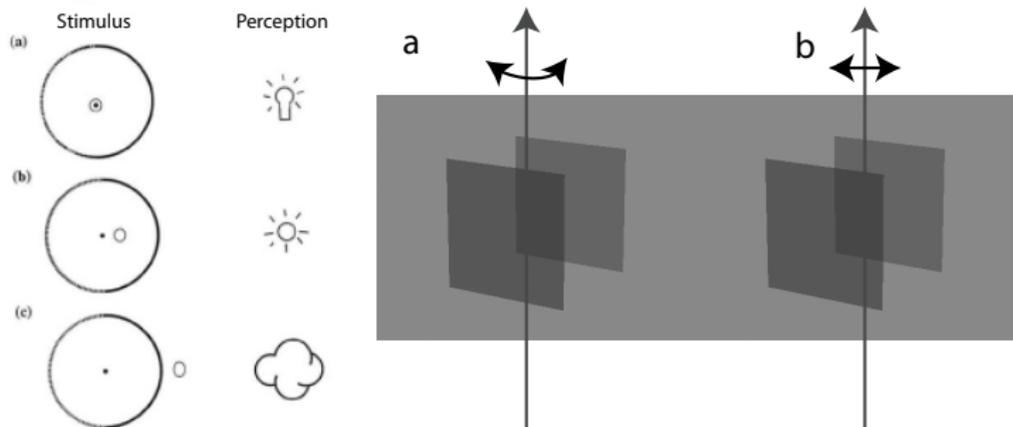


Figure 38: Examples of strong coupling with model selection. Left: A sphere is viewed binocularly, and small changes in the position of the specularity lead to very different percepts (Blake and Bülhoff, 1990). Right: Similarly altering the transparency of the moving surfaces can make the two surfaces appear to rotate either rigidly together or independently.

Model selection and explaining away

- ▶ Model selection can also give an alternative explanation for “explaining away”
- ▶ For example, consider two alternative models for partially occluded T
- ▶ The first model is of two individual segments plus a smudge region. The second is a T that is partially hidden by a smudge. The second model is more plausible since it would be very unlikely, an accidental viewpoint (or alignment), for the smudge to happen to cover the missing part of the T , unless it really did occlude it.
- ▶ A similar argument can be applied to the Kanizsa triangle. One interpretation is three circles partly occluded by a triangle. The other is three partial circles arranged so that the missing parts of the circles are aligned. The first interpretation is judged to be most probable.

Flying carpet revisited

- ▶ Like Kersten's ball-in-a-box experiments, the flying carpet illusion requires estimating the depth and orientation of the ground plane (i.e., the beach), segmenting and recognizing the woman and the towel she is standing on, detecting the shadow, and then using the shadow cues, which requires making some assumptions about the lighting, to estimate that the towel is hovering above the ground.
- ▶ This is a very complex way to combine all the cues in this image. Observe that it relies on the generic viewpoint assumption, in the sense that it is unlikely for there to be a shadow of that shape in that particular part of the image unless it was cast by some object. The real object that cast the shadow (the flag) is outside the image, so the visual system "attaches" the shadow to the towel, which then implies that the towel must be hovering off the ground.

Examples of strong coupling

We now give two examples of strong coupling. The first example deals with coupling different modalities, while the second example concerns the perception of texture.

Multisensory cue coupling

- ▶ Human observers are sensitive to both visual and auditory cues.
- ▶ Sometimes these cues have a common cause, e.g., you see a barking dog. But in other situations, the auditory and visual cues have different causes, e.g., a nearby cat moves and a dog barks in the distance.
- ▶ Ventriloquists are able to make the audience think that a puppet is talking by making it seem that visual cues (the movement of the puppet's head) and auditory cues (words spoken by the ventriloquist) are related. The ventriloquism effect occurs when visual and auditory cues have different causes – and so are in conflict – but the audience perceives them as having the same cause.

Multisensory cue coupling: The model (I)

- ▶ We describe an ideal observer for determining whether two cues have a common cause or not (Kording et al., 2007), which gives a good fit to experimental findings.
- ▶ The model is formulated using a meta-variable C , where $C = 1$ means that the cues x_A, x_V are coupled.
- ▶ More precisely, they are generated by the same process S by a distribution $P(x_A, x_V|S) = P(x_A|S)P(x_V|S)$.
 $P(x_A|S)$ and $P(x_V|S)$ are normal distributions $N(x_A|S, \sigma_A^2)$, $N(x_V|S, \sigma_V^2)$ – with the same mean S and variances σ_A^2, σ_V^2 .
- ▶ It is assumed that the visual cues are more precise than the auditory cues, so that $\sigma_A^2 > \sigma_V^2$. The true position S is drawn from a probability distribution $P(S)$, which is assumed to be a normal distribution $N(0, \sigma_p^2)$.

Multisensory cue coupling: The model (II)

- ▶ $C = 2$ means that the cues are generated by two different processes S_A and S_B .
- ▶ In this case, the cues x_A and x_V are generated respectively by $P(x_A|S_A)$ and $P(x_V|S_V)$, which are both Gaussian $N(S_A, \sigma_A^2)$ and $N(S_V, \sigma_V^2)$. We assume that S_A and S_V are independent samples from the normal distribution $N(0, \sigma_p^2)$.
- ▶ Note that this model involves model selection, between $C = 1$ and $C = 2$, and so, in vision terminology, it is a form of strong coupling with model selection (Yuille & Bulthoff, 1996).

Multisensory cue coupling: Illustration

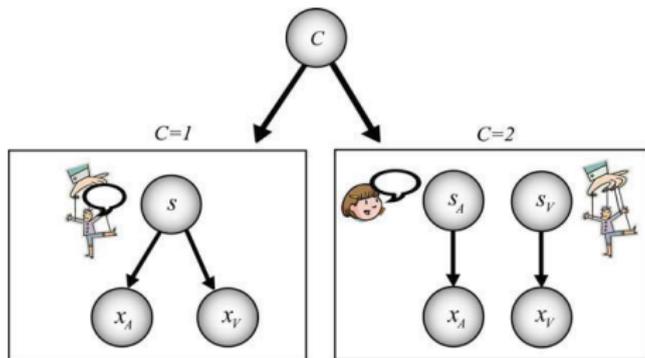


Figure 39: The subject is asked to estimate the position of the cues and to judge whether the cues are from a common cause – i.e., at the same location – or not. In Bayesian terms, the task of judging whether the cause is common can be formulated as model selection: are the auditory and visual cues more likely generated by a single cause (left) or by two independent causes (right)? Figure adapted from Kording et al. (2007).

Multisensory cue coupling: Comparison with experiments (I)

- ▶ This model was compared to experiments in which brief auditory and visual stimuli were presented simultaneously, with varying amounts of spatial disparity.
- ▶ Subjects were asked to identify the spatial location of the cue and/or whether they perceived a common cause (Wallace et al., 2004).
- ▶ The closer the visual stimulus was to the audio stimulus, the more likely subjects would perceive a common cause.
- ▶ In this case subjects' estimate of the stimuli's position was strongly biased by the visual stimulus (because it is considered more precise with $\sigma_V^2 < \sigma_A^2$).
- ▶ But if subjects perceived distinct causes, then their estimate was pushed away from the visual stimulus, and exhibited *negative bias*.

Multisensory cue coupling: Comparison with experiments (II)

- ▶ Körding et al. (2007) argue that this negative bias is a selection bias stemming from restricting to trials in which causes are perceived as being distinct.
- ▶ For example, if the auditory stimulus is at the center and the visual stimulus at 5 degrees to right of center, then sometimes the (very noisy) auditory cue will be close to the visual cue and hence judged to have a common cause, while in other cases, the auditory cause is farther away (more than 5 degrees).
- ▶ Hence the auditory cue will have a truncated Gaussian (if judged to be distinct) and will yield negative bias.

Multisensory cue coupling: Results and figure

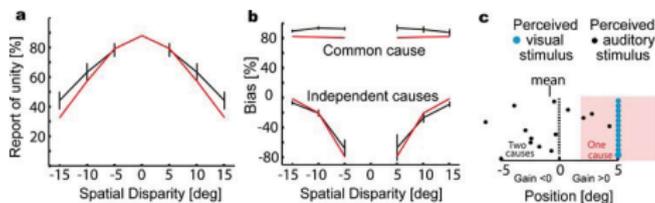


Figure 40: Reports of causal inference. (a) The relative frequency of subjects reporting one cause (black) is shown, with the prediction of the causal inference model (red). (b) The bias, i.e., the influence of vision on the perceived auditory position, is shown (gray and black). The predictions of the model are shown in red. (c) A schematic illustration explaining the finding of negative biases. Blue and black dots represent the perceived visual and auditory stimuli, respectively. In the pink area, people perceive a common cause. Reprinted with permission from Kording et al. (2007)

Multisensory cue coupling: The mathematics (I)

More formally, the beliefs $P(C|x_A, x_V)$ in these two hypotheses $C = 1, 2$ are obtained by summing out the estimated positions s_A, s_B of the two cues as follows:

$$\begin{aligned} P(C|x_A, x_V) &= \frac{P(x_A, x_V|C)P(C)}{P(x_A, x_V)} \\ &= \frac{\int dS P(x_A|S)P(x_V|S)P(S)}{P(x_A, x_V)}, \quad \text{if } C = 1, \\ &= \frac{\int \int dS_A dS_V P(x_A|S_A)P(x_V|S_V)P(S_A)P(S_V)}{P(x_A, x_V)}, \quad \text{if } C = 2. \end{aligned}$$

Multisensory cue coupling: The mathematics (II)

- ▶ There are two ways to combine the cues. The first is model selection. This estimates the most probable model $C^* = \arg \max P(C|x_V, x_A)$ from the input x_A, x_V and then uses this model to estimate the most likely positions s_A, s_V of the cues from the posterior distribution:

$$P(s_V, s_A) \approx P(s_V, s_A|x_V, x_A, C^*) = \frac{P(x_V, x_A|s_V, s_A, C^*)P(s_V, s_A|C^*)}{P(x_V, x_A|C^*)}.$$

- ▶ The second way to combine the cues is by *model averaging*. This does not commit itself to choosing C^* but instead averages over both models:

$$\begin{aligned} P(s_V, s_A|x_V, x_A) &= \sum_C P(s_V, s_A|x_V, x_A, C)P(C|x_V, x_A) \\ &= \sum_C \frac{P(x_V, x_A|s_V, s_A, C)P(s_V, s_A|C)P(C|x_V, x_A)}{P(x_V, x_A|C)}, \end{aligned}$$

where $P(C = 1|x_V, x_A) = \pi_C$ (the posterior mixing proportion).

Multisensory cue coupling: Extension

- ▶ Natarajan et al. (2008) showed that a variant of the model could fit the experiments even better.
- ▶ They replaced the Gaussian distributions with alternative distributions that are less sensitive to rare events. Gaussian distributions are non-robust because the tails of their distributions fall off rapidly, which gives very low probability to rare events.
- ▶ More precisely Natarajan et al. (2008) assumed that the data is distributed by a mixture of a Gaussian distribution, as above, and a uniform distribution (yielding longer tails).
- ▶ More formally, they assume $x_A \pi N(x_A : s_A, \sigma_A^2) + \frac{(1-\pi)}{r_1}$ and $x_V \pi N(x_V : s_V, \sigma_V^2) + \frac{(1-\pi)}{r_1}$, where π is a mixing proportion, and $U(x) = 1/r_1$ is a uniform distribution defined over the range r_1 .

Homogeneous and isotropic texture

- ▶ The second example is by Knill and concerns the estimating of orientation in depth (slant) from texture cues (Knill, 2003).
- ▶ There are alternative models for generating the image, and the human observer must infer which is most likely. In this example, the data could be generated by isotropic homogeneous texture or by homogeneous texture only.
- ▶ Knill's finding is that human vision is biased to interpret image texture as isotropic, but if enough data are available, the system turns off the isotropy assumption and interprets texture using the homogeneity assumption only.

Homogeneous and isotropic texture: Illustration

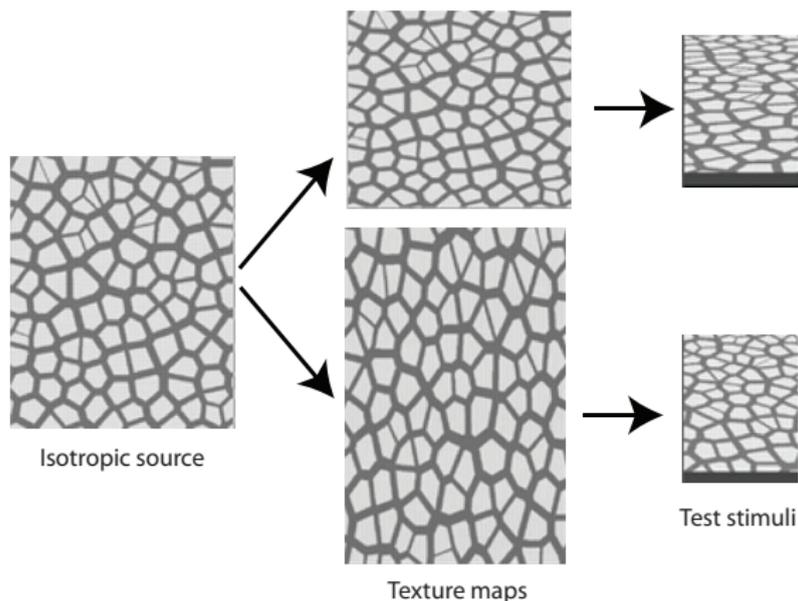


Figure 41: Generating textures that violate isotropy. An isotropic source image is either stretched (top middle) or compressed (bottom middle), producing texture maps that get applied to slanted surfaces shown on the right. A person that assumes surface textures are isotropic would overestimate the slant of the top stimulus and underestimate the slant of the bottom one. Figure adapted from from Knill (2003).

Homogeneous and isotropic texture: Theory (I)

- ▶ The posterior probability distribution for S is given by:

$$P(S|I) = \frac{P(I|S)P(S)}{P(I)}, \quad P(I|S) = \sum_{i=1}^n \phi_i P_i(I|S),$$

where ϕ_i is prior probability of model i , and $p_i(I|S)$ is corresponding likelihood function.

- ▶ More specifically, texture features T can be generated by either an isotropic surface or a homogeneous surface. The surface is parameterized by tilt and slant σ, τ . Homogeneous texture is described by two parameters α, θ , and isotropic texture is a special case where $\alpha = 1$. This gives two likelihood models for generating the data:

$$P_h(T|(\sigma, \tau), \alpha, \theta), \quad P_i(T|(\sigma, \tau), \theta)$$

Here, $P_i(T|(\sigma, \tau), \theta) = P_h(T|(\sigma, \tau), \alpha = 1, \theta)$.

Homogeneous and isotropic texture: Theory (II)

- ▶ Isotropic textures are a special case of homogeneous textures.
- ▶ The homogeneous model has more free parameters and hence has more flexibility to fit the data, which suggests that human observers should always prefer it. But the Occam factor (MacKay, 2003) means that this advantage will disappear if we put priors $P(\alpha)P(\theta)$ on the model parameters and integrate them out. This gives:

$$P_h(T|(\sigma, \tau)) = \int \int d\alpha d\theta P_h(T|(\sigma, \tau), \alpha, \theta),$$

$$P_i(T|(\sigma, \tau)) = \int d\theta P_h(T|(\sigma, \tau), \theta).$$

- ▶ Integrating over the model priors smooths out the models. The more flexible model, P_h , has only a fixed amount of probability to cover a large range of data (e.g., all homogeneous textures) and hence has lower probability for any specific data (e.g., isotropic textures).

Homogeneous and isotropic texture: The mathematics

- ▶ Knill describes how to combine these models using model averaging. The combined likelihood function is obtained by taking a weighted average:

$$P(T|(\sigma, \tau)) = p_h P_h(T|(\sigma, \tau)) + p_i P_i(T|(\sigma, \tau)), \quad (44)$$

where (p_h, p_i) are prior probabilities that the texture is homogeneous or isotropic. We use a prior $P(\sigma, \tau)$ on the surface and finally achieve a posterior:

$$P(\sigma, \tau|I) = \frac{P(I|(\sigma, \tau))P(\sigma, \tau)}{P(I)}. \quad (45)$$

- ▶ This model has a rich interpretation. If the data are consistent with an isotropic texture, then this model dominates the likelihood and strongly influences the perception. Alternatively, if the data are consistent only with homogeneous texture, then this model dominates. This gives a good fit to human performance (Knill, 2003).