Bayesian Decision Theory

Robert Jacobs
Department of Brain & Cognitive Sciences
University of Rochester

Types of Decisions

- Many different types of decision-making situations
  - Single decisions under uncertainty
    - Ex: Is a visual object an apple or an orange?
  - Sequences of decisions under uncertainty
    - Ex: What sequence of moves will allow me to win a chess game?
  - Choice between incommensurable commodities
    - Ex: Should we buy guns or butter?
  - Choices involving the relative values a person assigns to payoffs at different moments in time
    - Ex: Would I rather have $100 today or $105 tomorrow?
  - Decision making in social or group environments
    - Ex: How do my decisions depend on the actions of others?

Normative Versus Descriptive Decision Theory

- Normative: concerned with identifying the best decision to make assuming an ideal decision maker who is:
  - fully informed
  - able to compute with perfect accuracy
  - fully rational

- Descriptive: concerned with describing what people actually do

Decision Making Under Uncertainty

- Pascal's Wager:

	God exists	God does not exist
Live as if God exists	∞ (heaven)	0
Live as if God does not exist	$-\infty$ (hell)	0

- Expected payoff of believing in God is greater than the expected payoff of not believing in God
  - Believe in God!!!

Outline

- Signal Detection Theory
- Bayesian Decision Theory
- Dynamic Decision Making
  - Sequences of decisions

Signal Detection Theory (SDT)

- SDT used to analyze experimental data where the task is to categorize ambiguous stimuli which are either:
  - Generated by a known process (signal)
  - Obtained by chance (noise)

- <u>Example</u>: Radar operator must decide if radar screen indicates presence of enemy bomber or indicates noise

Signal Detection Theory

- Example: Face memory experiment
  - Stage 1: Subject memorizes faces in study set
  - Stage 2: Subject decides if each face in test set was seen during Stage 1 or is novel

- Decide based on internal feeling (sense of familiarity)
  - Strong sense: decide face was seen earlier (signal)
  - Weak sense: decide face was not seen earlier (noise)

Signal Detection Theory

	Decide Yes	Decide No
Signal Present	Hit	Miss
Signal Absent	False Alarm	Correct Rejection

- Four types of responses are not independent
  Ex: When signal is present, proportion of hits and proportion of misses sum to 1

Signal Detection Theory

- Explain responses via two parameters:
  - Sensitivity: measures difficulty of task
    - when task is easy, signal and noise are well separated
    - when task is hard, signal and noise overlap
  - Bias: measures strategy of subject
    - subject who always decides "yes" will never have any misses
    - subject who always decides "no" will never have any hits

- Historically, SDT is important because previous methods did not adequately distinguish between the real sensitivity of subjects and their (potential) response biases.

SDT Model Assumptions

- Subject’s responses depend on intensity of a hidden variable (e.g., familiarity of a face)
- Subject responds “yes” when intensity exceeds threshold
- Hidden variable values for noise have a Normal distribution
- Signal is added to the noise
  - Hidden variable values for signal have a Normal distribution with the same variance as the noise distribution

SDT Model

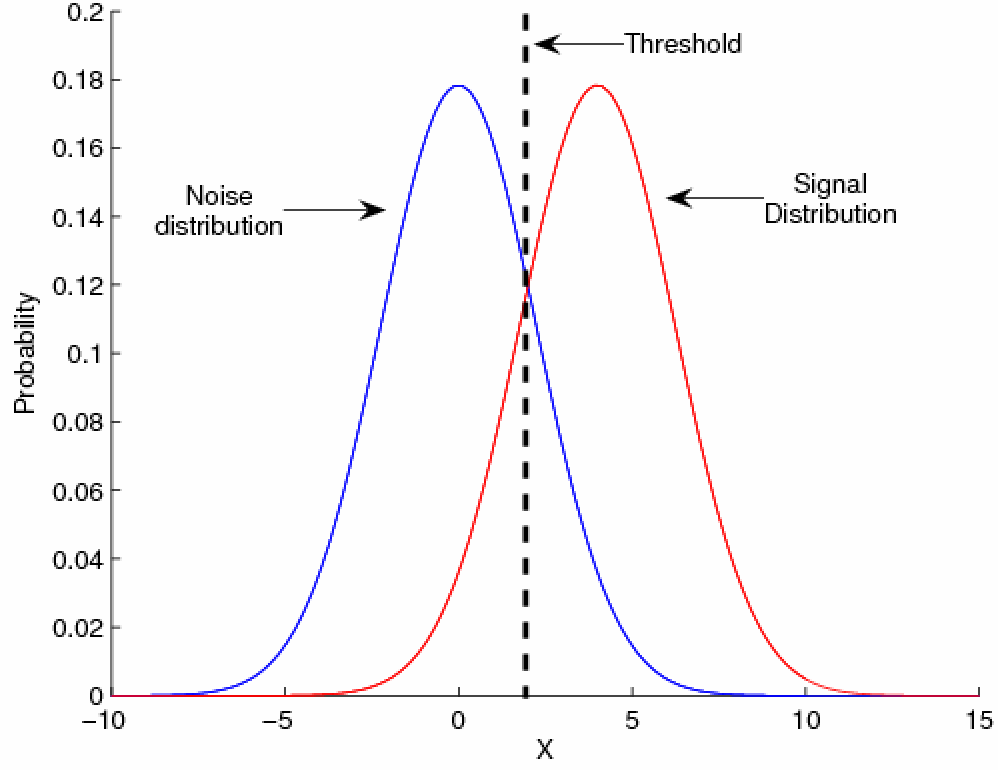

Threshold
Noise
distribution
Signal
Distribution
Probability
X
0.2
0.18
0.16
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
−10
−5
0
5
10
15

SDT Model

- Measure of sensitivity (independent of biases):

$$d'_{optimal} = \frac{\mu_s - \mu_N}{\sigma}$$

- Given assumptions, its possible to estimate $d'_{subject}$ from number of hits and false alarms

- Subject's efficiency:

$$\text{Efficiency} = \frac{d'_{subject}}{d'_{optimal}}$$

Bayesian Decision Theory

- Statistical approach quantifying tradeoffs between various decisions using probabilities and <u>costs</u> that accompany such decisions

- Example: Patient has trouble breathing
  - Decision: Asthma versus Lung cancer
  - Decide lung cancer when person has asthma
    - Cost: moderately high (e.g., order unnecessary tests, scare patient)
  - Decide asthma when person has lung cancer
    - Cost: very high (e.g., lose opportunity to treat cancer at early stage, death)

Example

- Fruits enter warehouse on a conveyer belt. Decide apple versus orange.

- w = type of fruit
  - w_1 = apple
  - w_2 = orange
- $P(w_1)$ = prior probability that next fruit is an apple
- $P(w_2)$ = prior probability that next fruit is an orange

Decision Rules

- Progression of decision rules:
  - (1) Decide based on prior probabilities
  - (2) Decide based on posterior probabilities
  - (3) Decide based on risk

(1) Decide Using Priors

- Based solely on prior information:

$$Decide \quad \begin{cases} w_1 & P(w_1) > P(w_2) \\ w_2 & otherwise \end{cases}$$

- What is probability of error?

$$P(error) = \min[P(w_1), P(w_2)]$$

(2) Decide Using Posteriors

- Collect data about individual item of fruit
  - Use lightness of fruit, denoted x, to improve decision making
- Use Bayes rule to combine data and prior information
- Class-Conditional probabilities
  - $p(x \mid w_1)$ = probability of lightness given apple
  - $p(x \mid w_2)$ = probability of lightness given orange

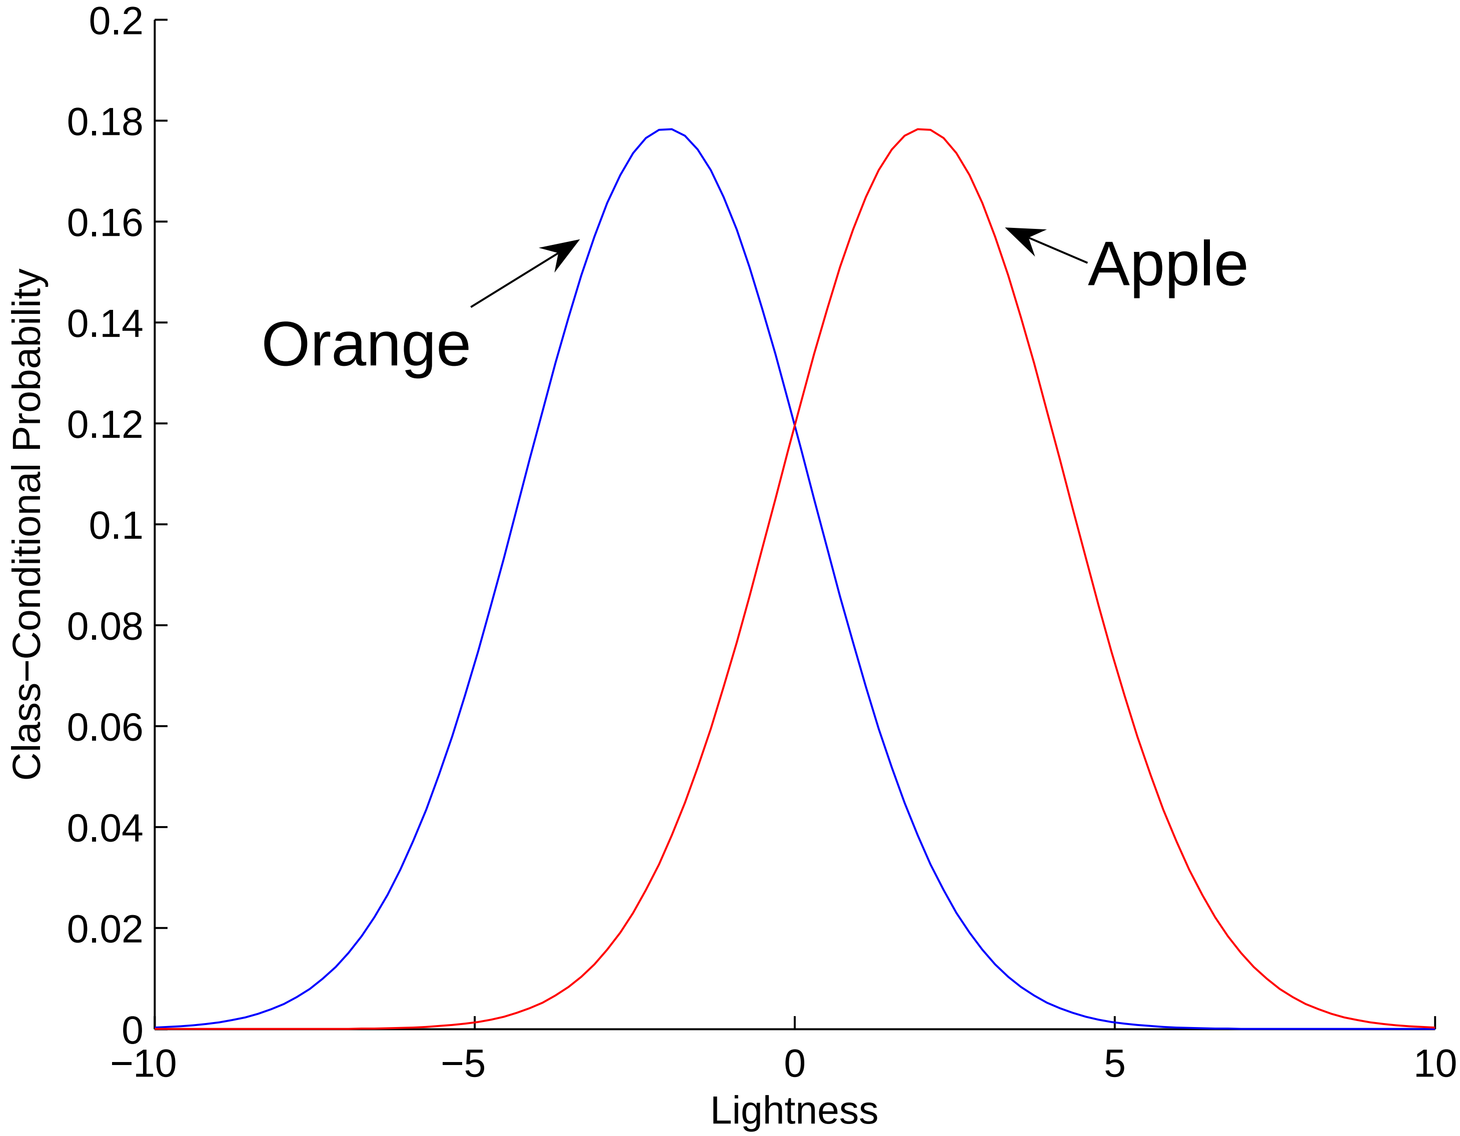
Class−Conditional Probability
0.2
0.18
0.16
0.14
0.12
0.1
0.08
0.06
0.04
0.02
0
−10
−5
0
5
10
Lightness
Orange
Apple

Bayes' Rule

- Posterior probabilities:

Likelihood

Prior

$$P(w_i \mid x) = \frac{p(x \mid w_i)\ p(w_i)}{p(x)}$$

Bayes Decision Rule

$$Decide \quad \begin{cases} w_1 & P(w_1 \mid x) > P(w_2 \mid x) \\ w_2 & otherwise \end{cases}$$

- Probability of error:

$$P(error \mid x) = \min[P(w_1 \mid x), P(w_2 \mid x)]$$

Assume equal prior probabilities:

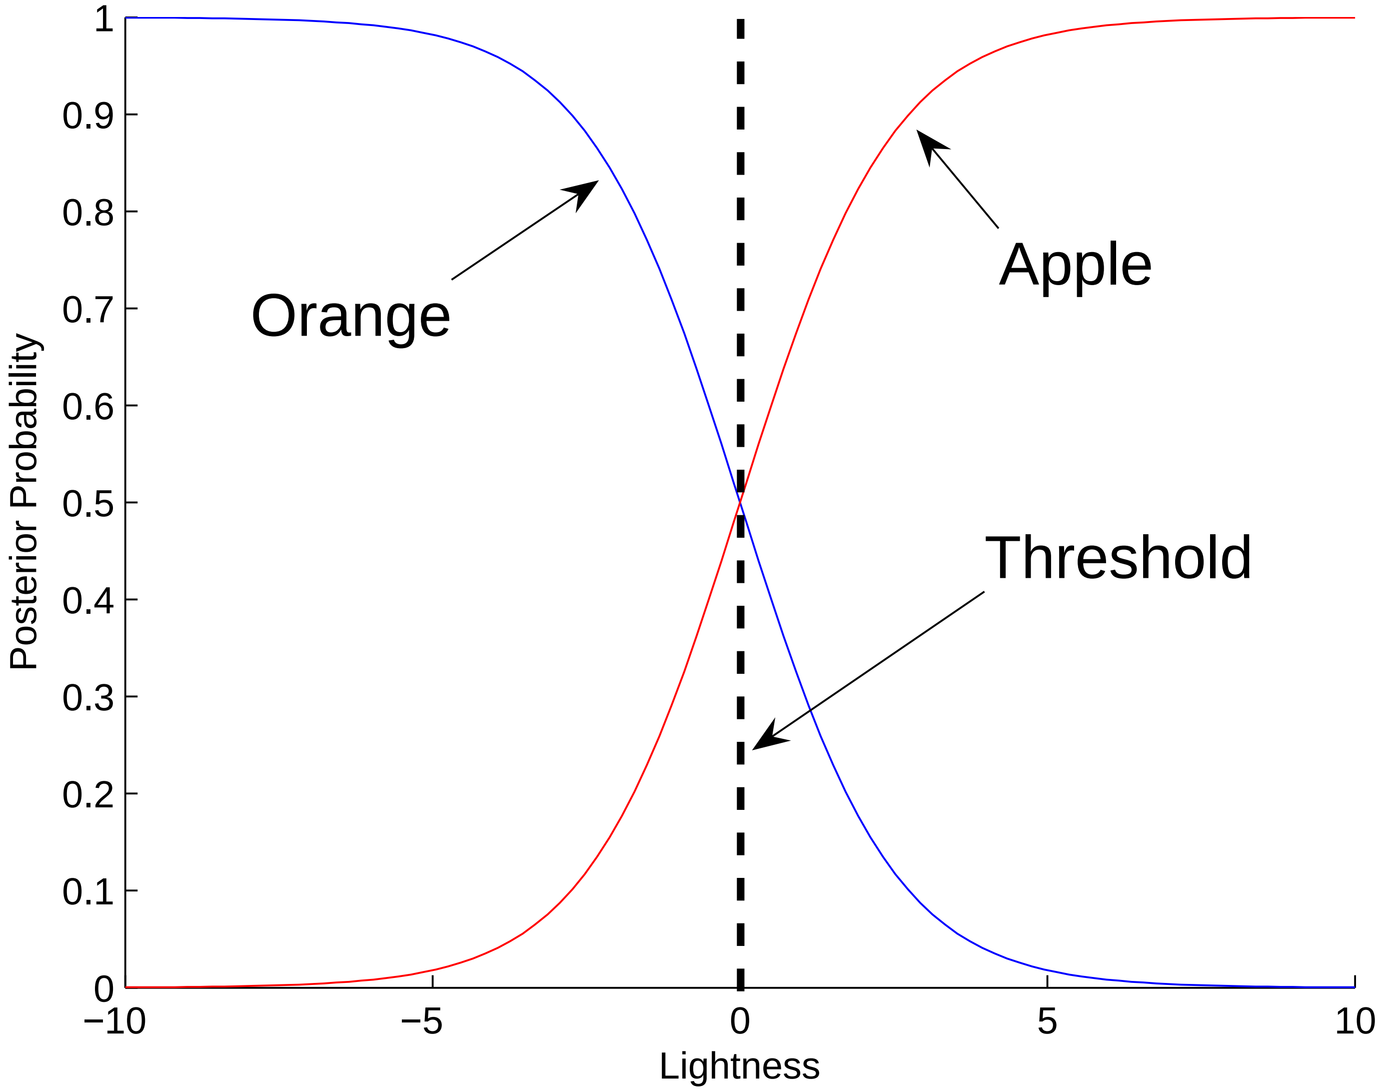

Prior probabilities: *P(orange) > P(apple)*

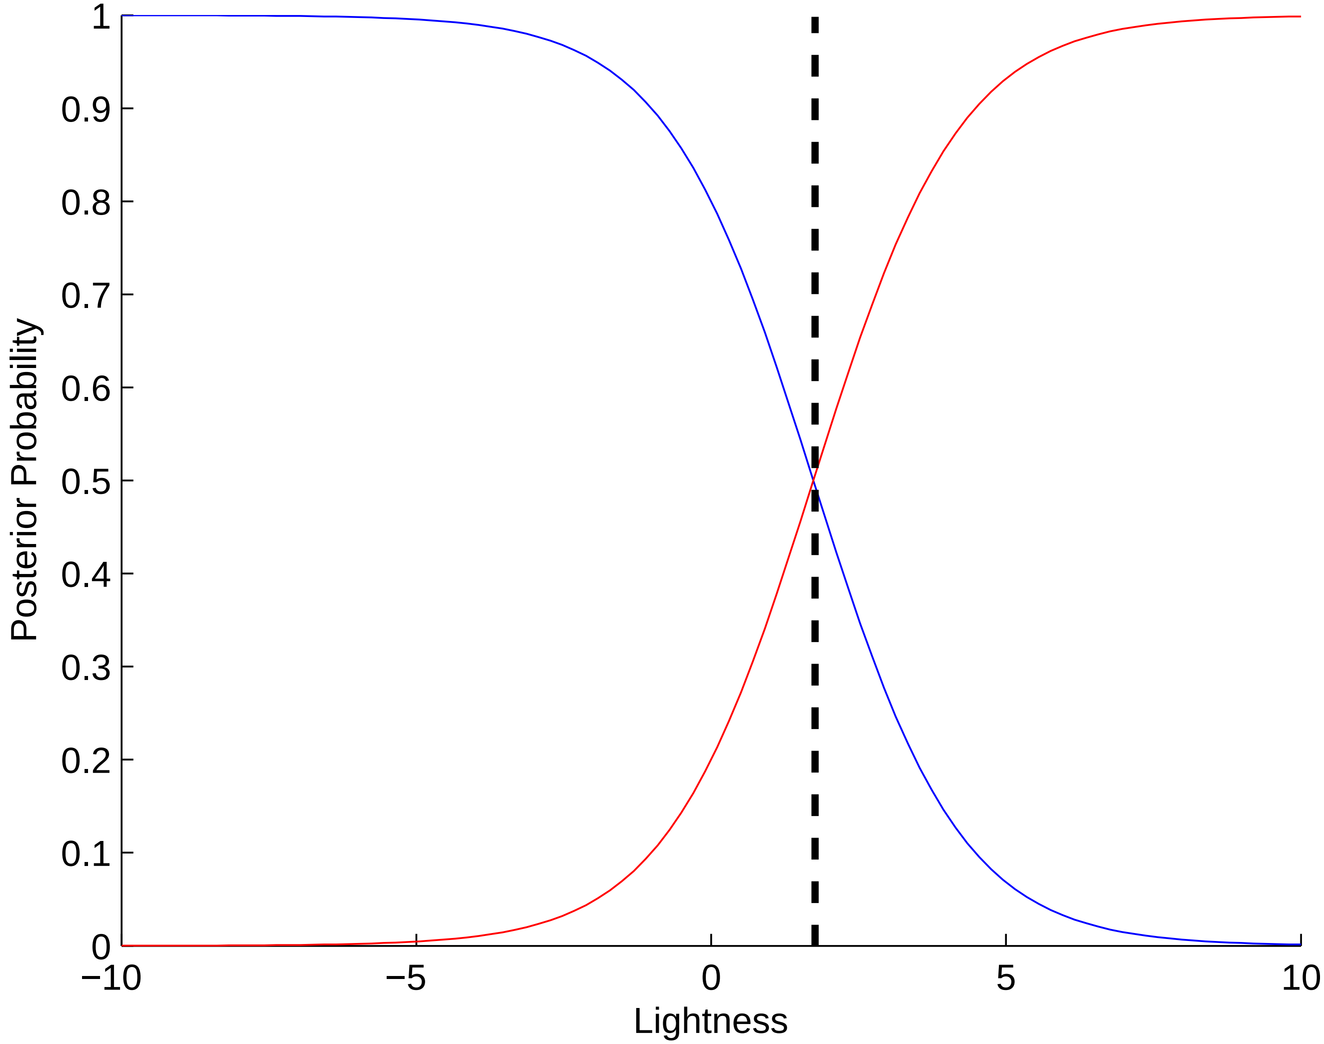

(3) Decide Using Risk

- $L(a_i \mid w_j)$ = loss incurred when take action a_i and the true state of the world is w_j

- Expected loss (or conditional risk) when taking action a_i:

$$R(a_i \mid x) = \sum_j L(a_i \mid w_j)\, P(w_j \mid x)$$

Loss function   Posterior

Minimum Risk Classification

- $a(x)$ = decision rule for choosing an action when x is observed

- Bayes decision rule: minimize risk by selecting the action a_i for which $R(a_i \mid x)$ is minimum

Loss Functions for Classification

- Zero-One Loss
  - If decision correct, loss is zero
  - If decision incorrect, loss is one
- What if we use an asymmetric loss function?
  - $L(apple \mid orange) > L(orange \mid apple)$

$L(apple \mid orange) > L(orange \mid apple)$

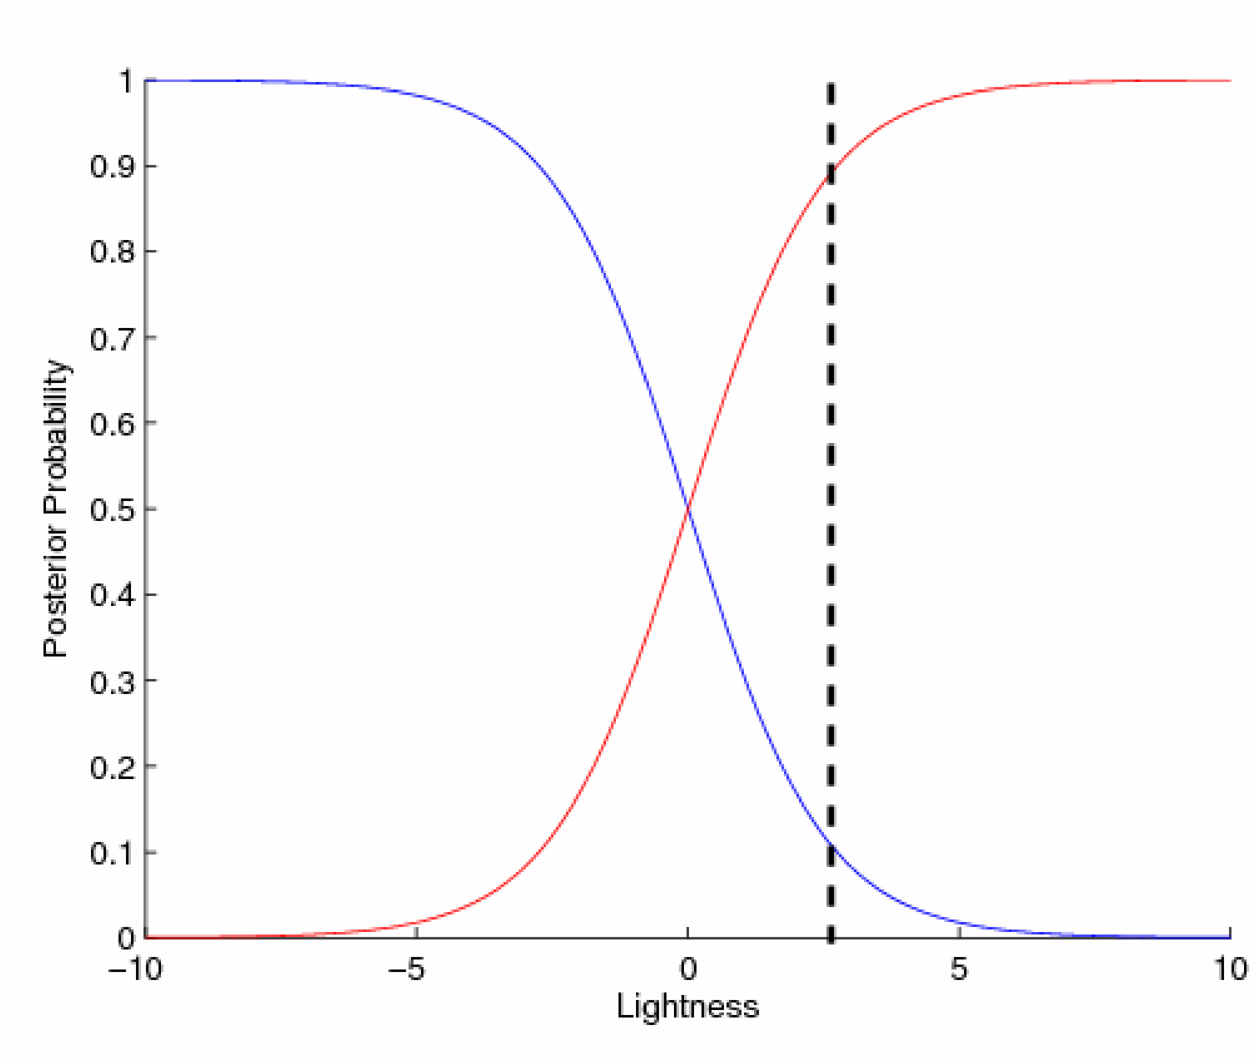

Loss Functions for Regression

- Delta function
  - $L(y|y^*) = -\delta(y-y^*)$
  - Optimal decision: MAP estimate
    - action y that maximizes $p(y \mid x)$ [i.e., mode of posterior]

- Squared Error

$$L(y \mid y^*) = (y - y^*)^2$$

  - Optimal decision: mean of posterior

Loss Functions for Regression

- Local Mass Loss Function

$$L(y \mid y^{*}) = -\exp[-\frac{(y - y^{*})^{2}}{\sigma^{2}}]$$

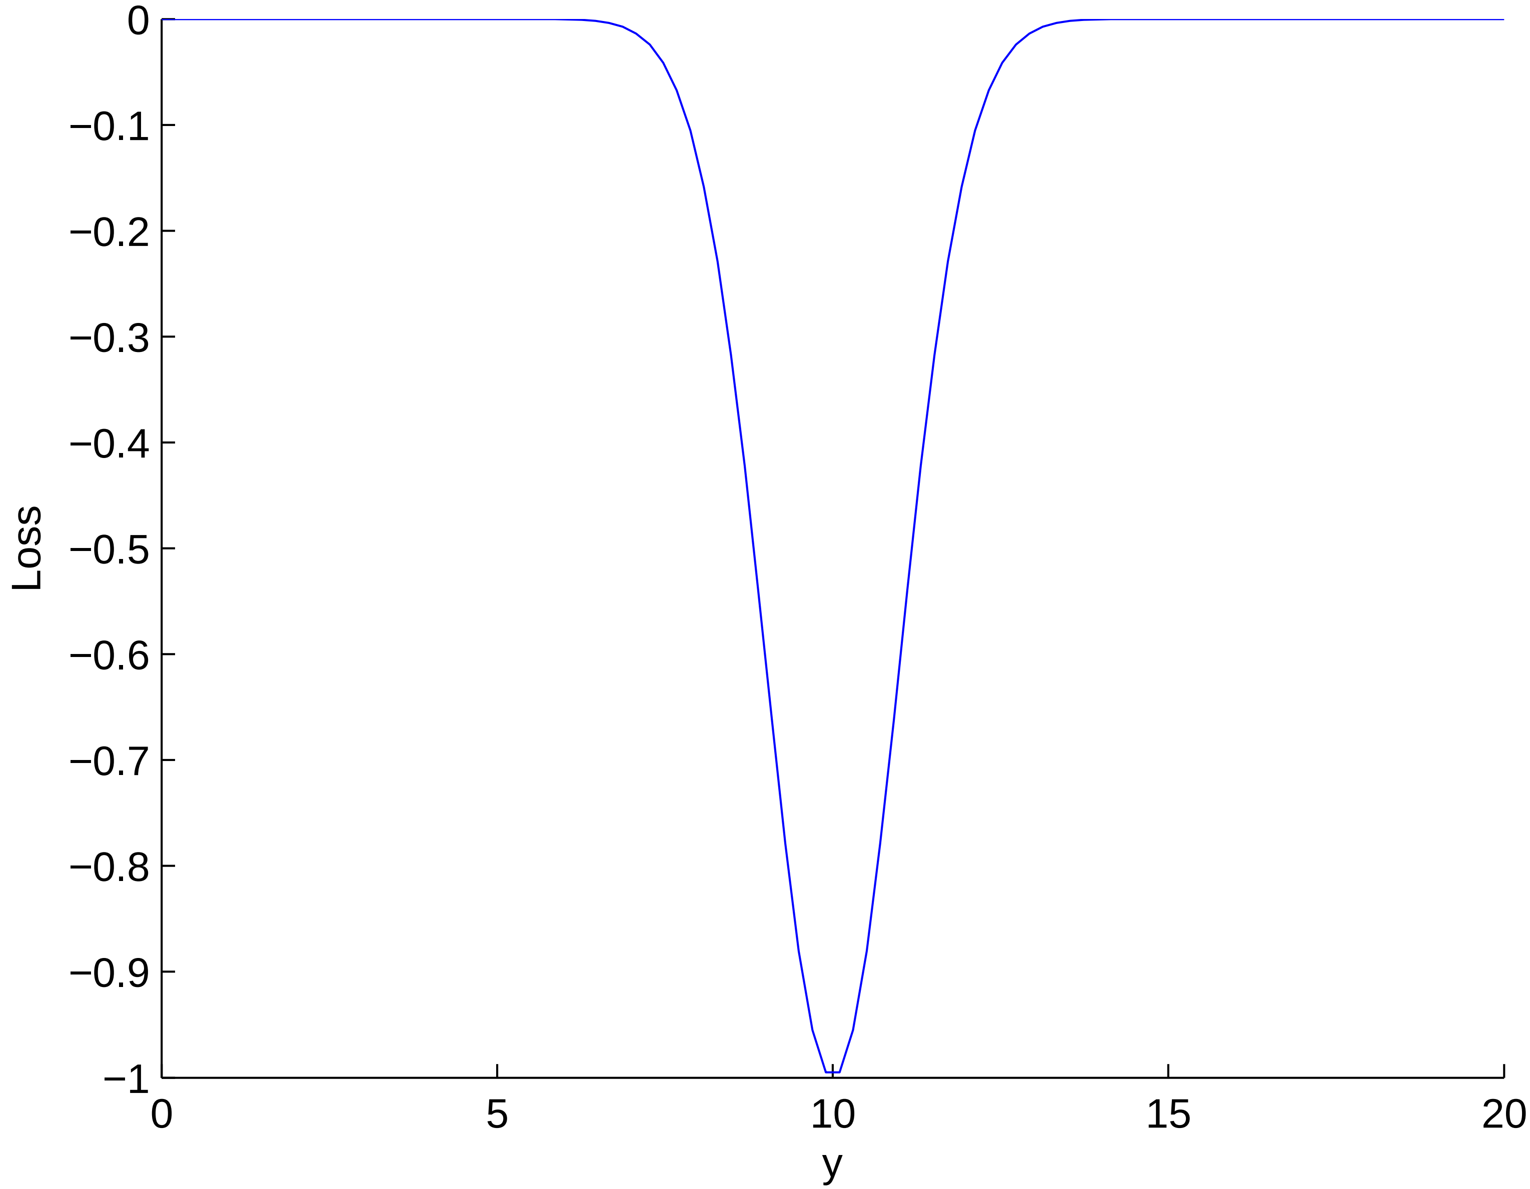

0
−0.1
−0.2
−0.3
−0.4
−0.5
−0.6
−0.7
−0.8
−0.9
−1
Loss
0
5
10
15
20
y

Freeman (1996): Shape-From-Shading

- Problem: Image is compatible with many different scene interpretations
- Solution: Generic view assumption
  - Scene is not viewed from a special position

Generic Viewpoint Assumption

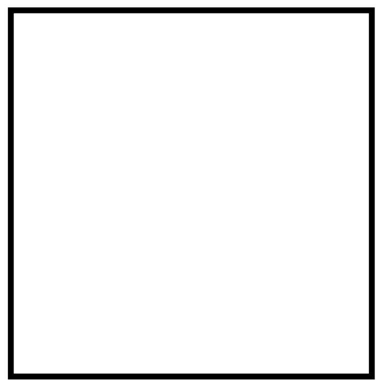

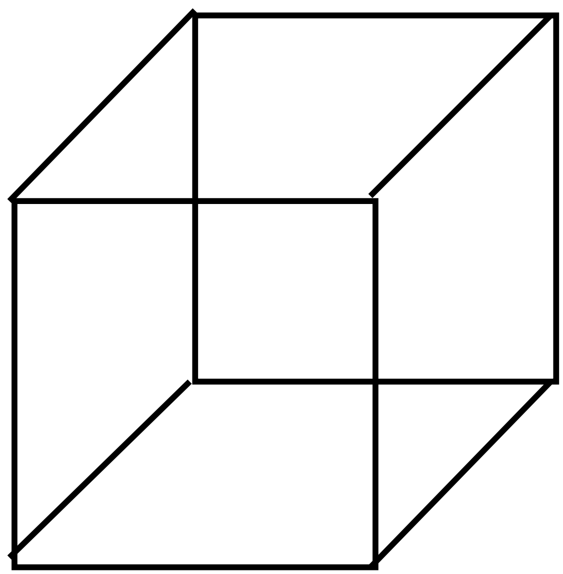

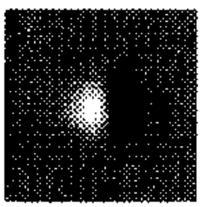

(a)

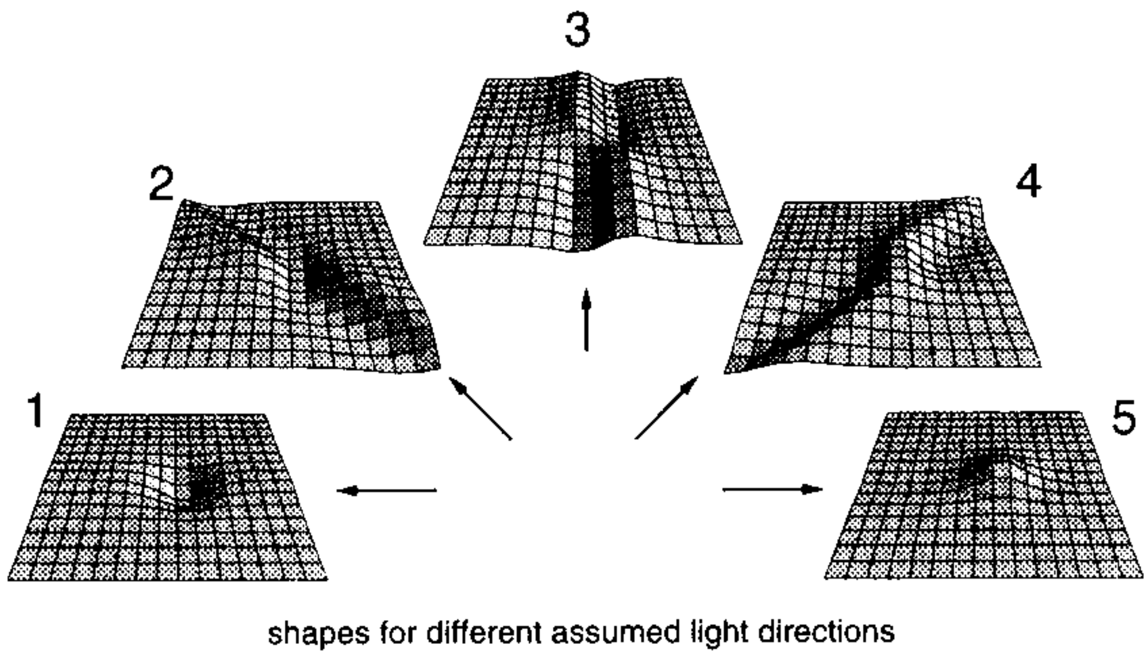

(b)

Figure from Freeman (1996)

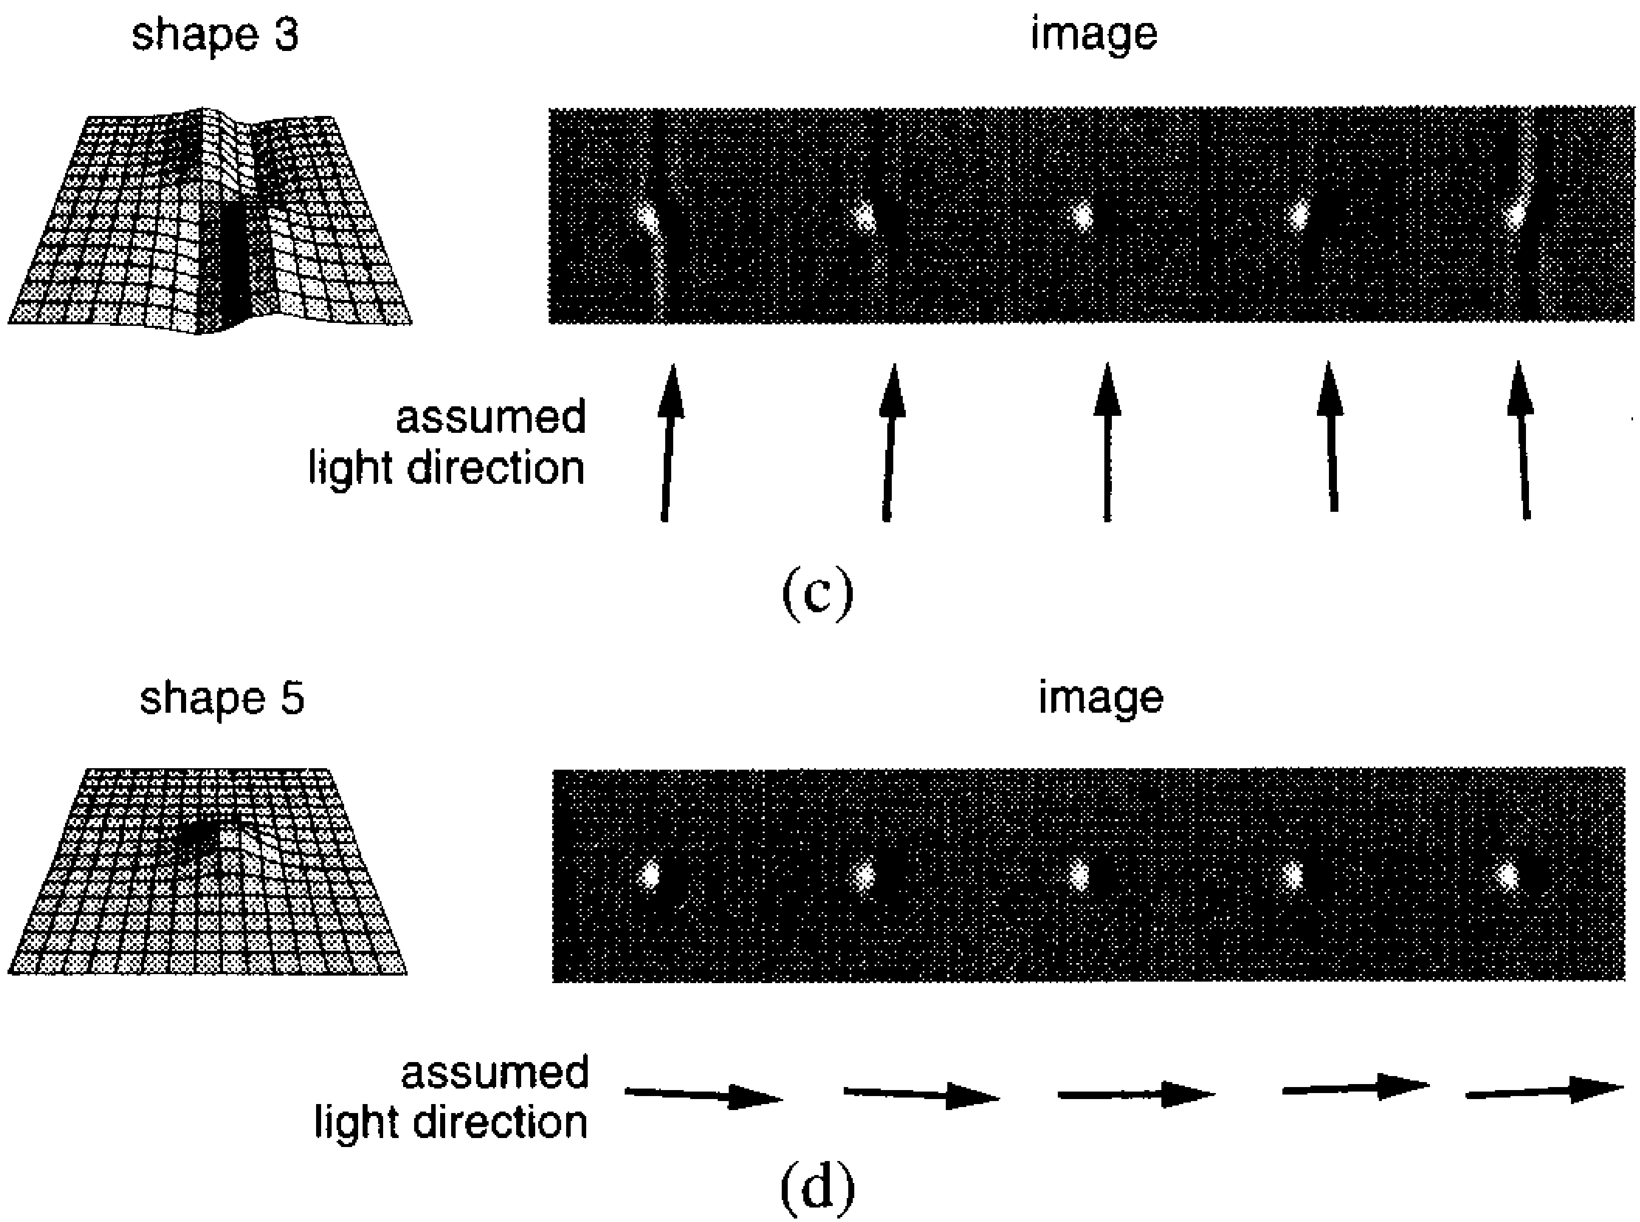

Figure from Freeman (1996)

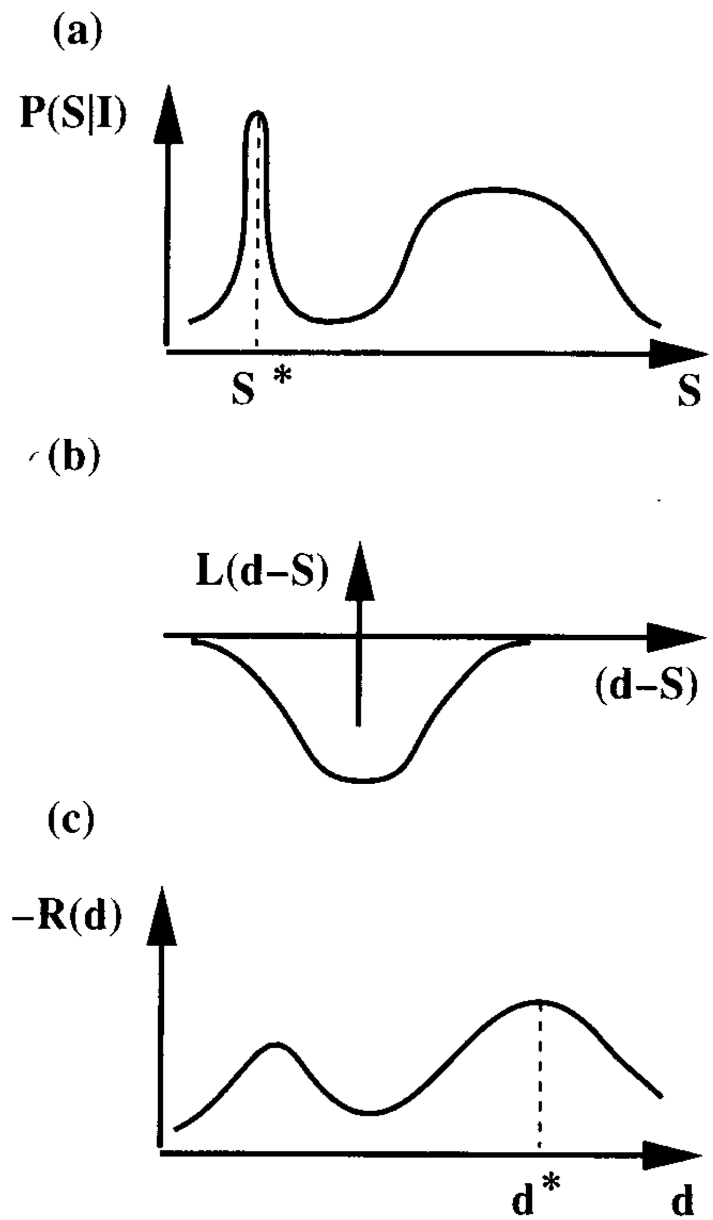

Figure from Yuille and Bülthoff (1996)

- Likelihood favors (b)
- Prior assuming smoothness favors (b)
- Loss function penalizing precise alignment between light source and object favors (c)

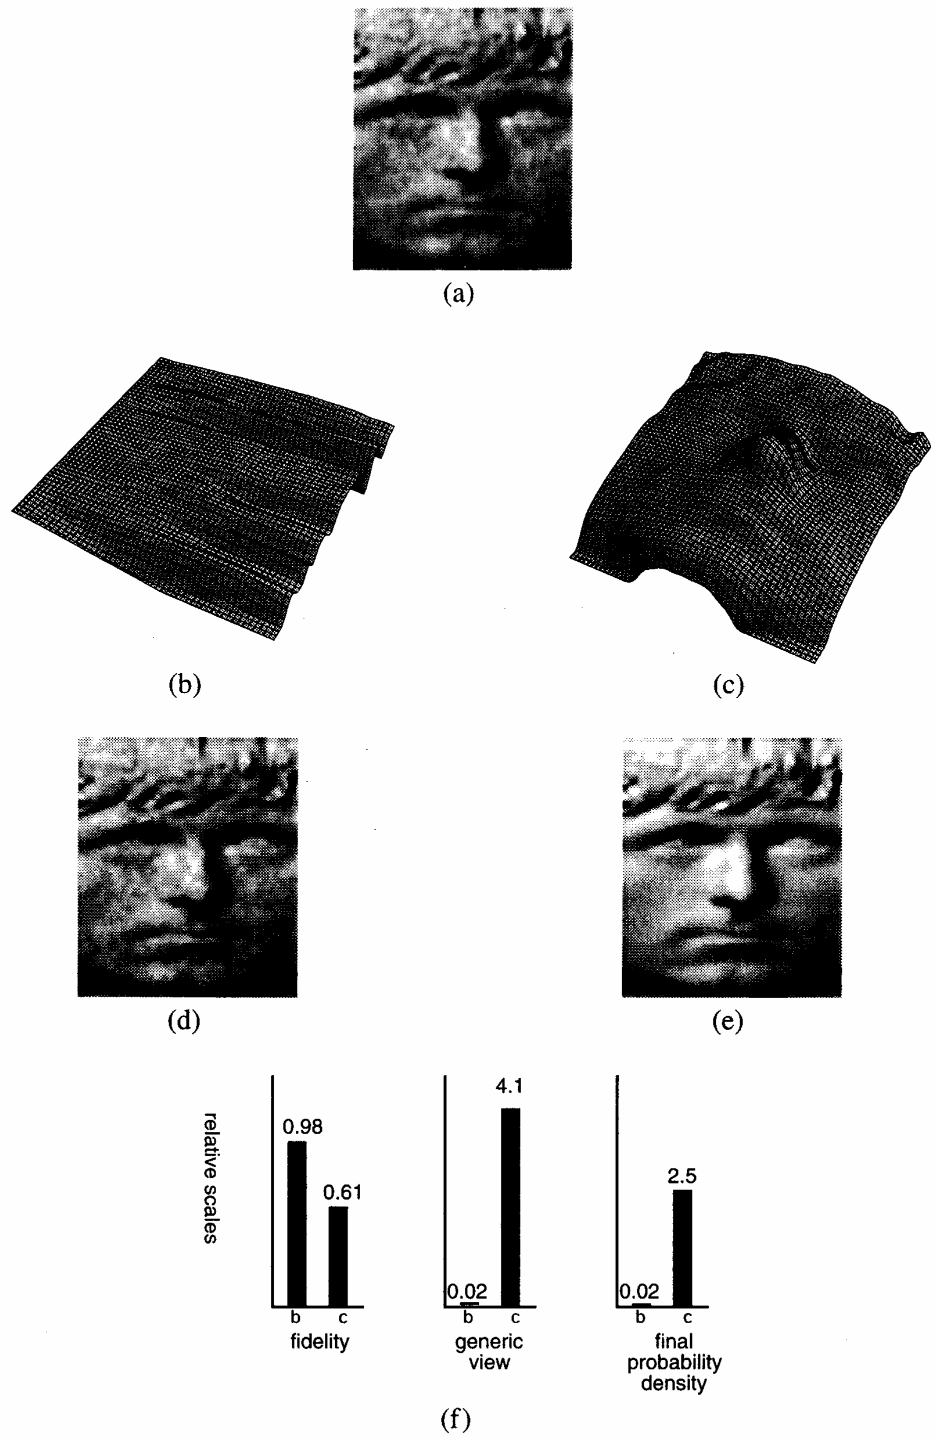

Figure from Freeman (1996)

Dynamic Decision Making

- Decision-making in environments with complex temporal dynamics
  - Decision-making at many moments in time
  - Temporal dependencies among decisions

- Examples:
  - Flying an airplane
  - Piloting a boat
  - Controlling an industrial process
  - Coordinating firefighters to fight a fire

Loss Function

- Example: Reaching task
  - Move finger from location *A* to location *B* within 350 msec

- Loss function
  - Finger should be near location *B* at end of movement
  - Velocity at end of movement should be zero
  - Movement should use a small amount of energy

- This loss function tends to produce smooth, straight motions

Markov Decision Processes (MDP)

- S is the state space
- A is the action space
- (State, Action)$_t$ → (State)$_{t+1}$
  - $P_a(s, s') = P(s_{t+1} = s' \mid s_t = s, a_t = a)$
- $R(s)$ = immediate reward received in state s

- Goal: choose actions so as to maximize discounted sum of future rewards

$$\sum_{t=0}^{\infty} \gamma^t R(s_t) \text{ with } 0 < \gamma \leq 1$$

Markov Decision Processes (MDP)

- Policy: mapping from states to actions
- Optimal policies can be found via dynamic programming
  - Caveat: computationally expensive!!!
  - Reinforcement learning: approximate dynamic programming

Summary

- Many different types of decision making situations
- Normative versus descriptive theories
- Signal detection theory
  - Measures sensitivity and bias independently
- Bayesian decision theory: single decisions
  - Decide based on priors
  - Decide based on posteriors
  - Decide based on risk
- Bayesian decision theory: dynamic decision making
  - Markov decision processes