


Bayesian reasoning in computer vision


Bill Freeman, Antonio Torralba, Phillip Isola

ambiguity in computer vision




<https://www.michaelbach.de/ot/sze-Necker/index.html>

ambiguity in computer vision



<http://lupuvictor.blogspot.com/2011/08/copernicus-crater-optical-illusion.html>

ambiguity in computer vision



size ambiguity



size ambiguity



Distribution of surface reflectances



0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09	0.10
0.11	0.12	0.13	0.14	0.15	0.16	0.17	0.18	0.19	0.20
0.21	0.22	0.23	0.24	0.25	0.26	0.27	0.28	0.29	0.30
0.31	0.32	0.33	0.34	0.35	0.36	0.37	0.38	0.39	0.40
0.41	0.42	0.43	0.44	0.45	0.46	0.47	0.48	0.49	0.50
0.51	0.52	0.53	0.54	0.55	0.56	0.57	0.58	0.59	0.60
0.61	0.62	0.63	0.64	0.65	0.66	0.67	0.68	0.69	0.70
0.71	0.72	0.73	0.74	0.75	0.76	0.77	0.78	0.79	0.80
0.81	0.82	0.83	0.84	0.85	0.86	0.87	0.88	0.89	0.90
0.91	0.92	0.93	0.94	0.95	0.96	0.97	0.98	0.99	1.00

Distribution of illuminant intensities

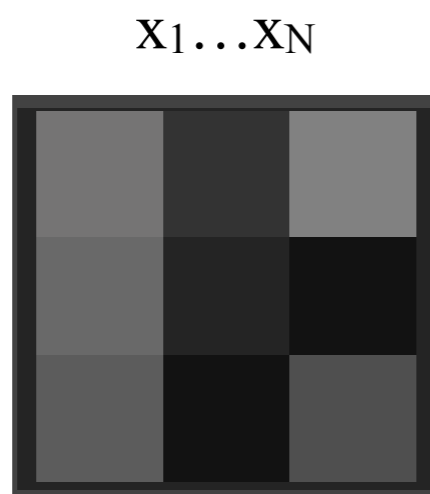
1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50
51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70
71	72	73	74	75	76	77	78	79	80
81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100

Surface brightness ambiguity


Observation, y



Bright light




Dim squares



Dim light

Bright squares

Likelihood term for a $b = 1$ problem



Role of priors in computer vision



Sinha and Poggio, 1996

Role of priors in computer vision



<https://www.michaelbach.de/ot/>

Role of the loss function


Picking a single best \mathbf{x}

From the supplementary notes for this lecture:


with the *loss function*, which specifies the penalty for guessing wrong. Let $\hat{\mathbf{x}}$ be your estimate of the parameters, \mathbf{x} . Then $L(\hat{\mathbf{x}}, \mathbf{x})$ is the loss incurred by guessing $\hat{\mathbf{x}}$ when the true value was \mathbf{x} . With the posterior probability, we can calculate the expected loss, $\bar{L}(\hat{\mathbf{x}}, \mathbf{x})$

$$\bar{L}(\hat{\mathbf{x}}, \mathbf{x}) = \int_{\mathbf{x}} L(\hat{\mathbf{x}}, \mathbf{x}) P(\mathbf{x}|\mathbf{y}) \quad (6.20)$$

We often use a loss function which is only a function of $\hat{\mathbf{x}} - \mathbf{x}$.




(a) Posterior Probability




(b) Ridge Thickness Variations


Two loss functions (left), and the (minus) expected losses for the $1=ab$ problem




(a) MAP loss fn.



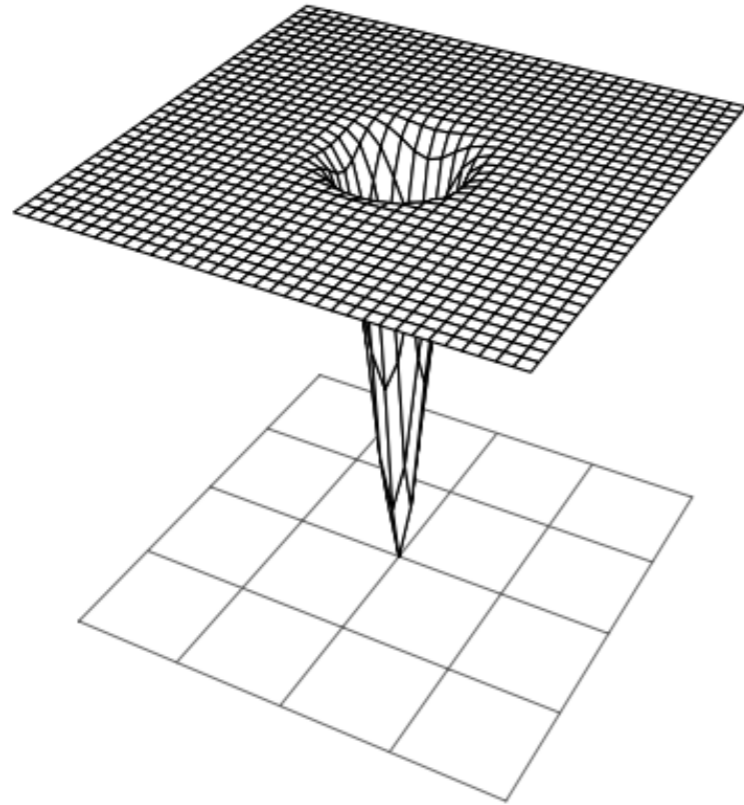
(d) (minus) MAP risk




(b) MMSE loss fn.



(e) (minus) MMSE risk



(c) MLM loss fn.



(f) (minus) MLM risk