

Probabilistic Models of the Visual Cortex

Fall 2021 Homework 3

Prof. Alan Yuille

Due on Nov 16 before the class. Late homework will not be accepted unless permission is obtained well in advance in documented extenuating circumstances. Submit pdf file on the Gradescope (entry code: posted in Piazza). Format file name as *firstname-lastname-hw3.pdf*. Do not include the iPython notebook code in the pdf submission as it is not required. Please start early on the coding part and make sure the software (i.e. the jupyter notebook) works for you. If you have any questions about the homework, please post on the Piazza Q&A - hw3.

Question 1. Cue Coupling (15 points)

What is the distinction between weak and strong coupling of visual cues? Under what conditions can cues be coupled by weighted linear summation? When does this show that cue coupling is optimal in the sense of combining cues taking into account their uncertainty? What is double counting, and how can it be avoided? Briefly discuss directed graphical models and how they can be applied to model the problem where your friend might have psychic powers. Give two examples of strong coupling. Briefly describe the multisensory coupling of audio and visual cues,

describing how this can be formulated as model selection to determine common cause (or not), and explaining why this formulation can exhibit negative bias.

Question 2. Gibbs Sampling (13 points)

The Ising model is specified by a Gibbs Distribution $P(\vec{S}|\vec{I}) = \frac{1}{Z} \exp\{-E(\vec{S}; \vec{I})\}$ where the energy $E(\vec{S}; \vec{I})$ can be expressed by:

$$E(\vec{S}; \vec{I}) = \sum_x (S(x) - I(x))^2 + \lambda \sum_x \sum_{y \in Nbh(x)} (S(x) - S(y))^2.$$

Here, $Nbh(x)$ denotes the set of pixel indices neighboring x , $S(x) \in \{0, 1\}$ (also called the *state*), and $I(x) \in [0, 1]$ (the *image*).

1. Describe how this model captures spatial context. (2points) What is the likelihood and prior for this model? (2points each)
2. Why is it impossible to sample directly from the full joint distribution $P(\vec{S}|\vec{I})$? (2 points). Compare to sample from the full joint distribution as above, what distribution is sampled from the Gibbs sampling? (2points) Derive the Gibbs sampling distribution for this model. (4 points) What theoretical results guarantee that Gibbs sampling will converge to samples from the Gibbs distribution (do not need details or derivations, just broad results). (3 points)

Question 3. Mean Field Theory (15 points)

1. Describe the mean field theory approximation for the Ising model $P(\vec{S}|\vec{I})$. (6 points) What is $Q(\vec{S})$ and how is the Kullback-Leiber divergence used as a measure of similarity between $P(\cdot)$ and $Q(\cdot)$? (6 points)
2. In mean field theory, what is the free energy? (3 points)

Question 4. Boltzmann Machine (12 points)

1. What is the difference between a Boltzmann Machine and a Restricted Boltzmann Machine? (3 points) What is the difference between the hidden and the observed variables (\vec{S}_o, \vec{S}_h) . $P(\vec{S}) = \frac{1}{Z} \exp\{-E(\vec{S})\}$, where $E(\vec{S}) = C \sum_{ij} \omega_{ij} S_i S_j$. (4 points) What is the update rule for learning the weights $\{\omega_{ij}\}$ in terms of the expected statistics of $S_i S_j$ with respect to the clamped and unclamped distributions. (3 points)
2. Why is it easier to learn these weights for the Restricted Boltzmann Machine. (2 points)

Question 5. Motion and Kalman Filtering (17 points)

1. List one difference between short-range and long-range motion. (1 point) What is one problem of estimating short-range motion? (2 points)
2. To estimate short-range motion, what are the priors terms and their equations? (4 points) Given the prior, if there is no observation, what are the estimated two dimensional velocities? (2 points) What are the estimated velocities if you have an observation term? (2 points)
3. What are the two stages of Bayes-Kalman filter and their equations? (4 points) What is the purpose of the second step? (1 point) Derive the second step equation by using Bayes' rule. (1 point)

Question 6. Experimental Section: foreground-background segmentation (16 points)

In this question, you will use Gibbs sampling to apply the Ising model to foreground-background segmentation, IPython notebook is used for this project, follow:

<https://github.com/ccvl/VisualCortexCourse/blob/master/HW3/Gibbs%20Sampling>.

ipynb

Note: The modeling conventions in the notebook are slightly different from those above. In particular, in the code, S takes values ± 1 and not $\{0, 1\}$.