

**1. User-specific Facebook advertisements.** Come up with some model that will help advertisers choose which Facebook users to show their ads. For example, you could model the likelihood that a certain Facebook user would click on a certain ad at a certain time. You may want to try to model the user's hidden state- e.g. hunger or happiness level might affect whether they would click.

**2. Predicting election outcomes.** Develop a model to predict presidential election outcomes in each of the 50 states. You may want to consider the simplified case where there are only two candidates. Think about underlying factors that might affect voting across all states, and about factors that might affect voting in a state-specific way. Could you provide some sort of confidence interval or estimate of uncertainty about your predictions? Maybe you would want to consider spatial relationships, such as the co-influence of neighboring states. It might be interesting to think about the time dimension as well.

**3. Removing noise from speech recordings.** Think about how you could filter out noise from human speech recordings that have additional 'non-Gaussian' noise such as birds chirping, babies crying, cars honking, etc. You could either consider the case where you know all of the different types of noise that you will have to deal with (and so perhaps you would want to explicitly model what bird chirping sounds like, etc.), or the case where you are expected to deal with any kind of noise (in which case you would probably just want to model human speech).

**4. Predicting traffic patterns.** Suppose you are given a map of interconnected roads. Come up with a model that could predict how heavy the traffic will be on a given road (or stretch of road) at a given time  $t$ , where you observe traffic load all the way up to time  $(t-1)$ . Try to provide some estimates of uncertainty.

**5. Modeling human movement.** You are hired by Pixar to come up with a generative model of different styles of dance, so that they can easily generate animated scenes of crowds of people dancing. You might want to think about using online videos (e.g. YouTube), which have titles, video, and audio. You can assume that you have video processing capabilities that can identify specific markers on individual bodies in the frame and infer 3-dimensional coordinates of the markers.

**6. Music in the digital age:** Think about how you could use tools from this class to address the following question: Does user sharing and/or copying music positively or negatively (or not) affect the success (e.g. monetary or degree of popularity) of recording artists? How could you measure and then model these different activities and outcomes?