Lecture Notes: Introduction to Filtering

Imagine you are using a radar to track a plane, and you want to shoot it down. You might see something like Figure 1. In deciding where to shoot, you face some problems. First, the radar measurements have noise. This means that each measurement has some error with respect to the actual position of the plane. After taking a lot of measurements, we can develop an understanding of the typical amount of expected error. Visually this can be seen in Figure 1. If you only look at the first 2 or 3 measurements, the amount of noise is not clear. But after looking at the entire figure, you can easily develop an intuitive feeling about the typical amount of noise the measurements. This information can be useful in deciding how to shoot. For example, knowing a reasonable limit on the bounds of the noise, we can pick a missile with a payload that will explode in a large enough radius to cover most of the noise.

The second problem is that the radar measurements are not continuous. For example, a common radar rotates a fan-shaped beam in 360 degrees about once every 10 seconds. Between consecutive measurements, we have no real data about the location of the plane. Instead, we must interpolate or fill in the expected location. Looking at Figure 1, you can visually see that the plane appears to be flying in a straight line. Based on this conceptual model, one can fill in the likely location of the plane between measurements. Figure 2 shows a likely model. This information can also be useful in deciding how to shoot. For example, if you need to fire a missile between radar scans, you can project ahead the expected location of the plane, and aim the missile to that location.

The third problem with measurements is that the variable of interest is not always directly observable. For the plane shooting problem, the plane's location is directly measured by the radar. But for other problems, the variable of interest is only indirectly observed through measurements of related variables. For example, consider the weather map in Figure 3. It is commonly used to show areas of weather probabilities, such as the likelihood of rain versus sunshine tomorrow. A weather map is also commonly used to show cold and warm fronts, and their likely effect on the weather over the next few days. These variables of interest are not directly observed. Instead, variables like temperature, barometric pressure, and humidity are observed, and from these, the variables of interest are inferred. Just like we used a model to fill in the plane's position between measurements, we can use a model to relate the measured variables to the variables of interest.

This course is concerned with using mathematical models of systems and observations to improve tracking. This technique is called filtering. The models are equations that describe expected or observed variables of interest. The tracking problem could be anything from tracking the physical location of an object, to tracking weather or economic probabilities, to tracking broadcast signals, to tracking the health or physiology of a person. The

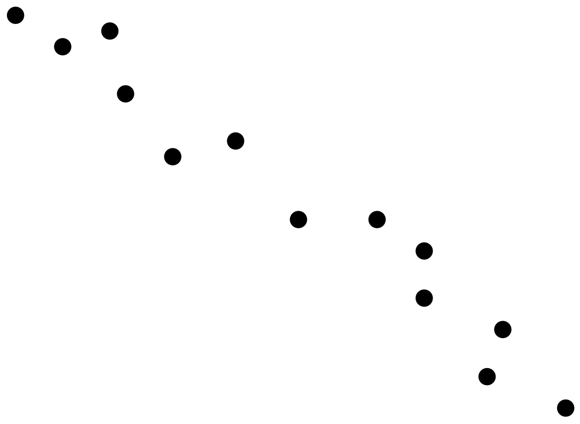

Figure 1: Radar measurements (positions) of a plane moving from the upper-left towards the lower-right.

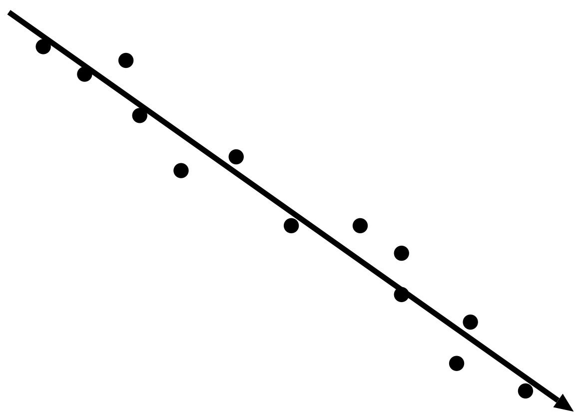

Figure 2: Linear model of the motion of the plane.

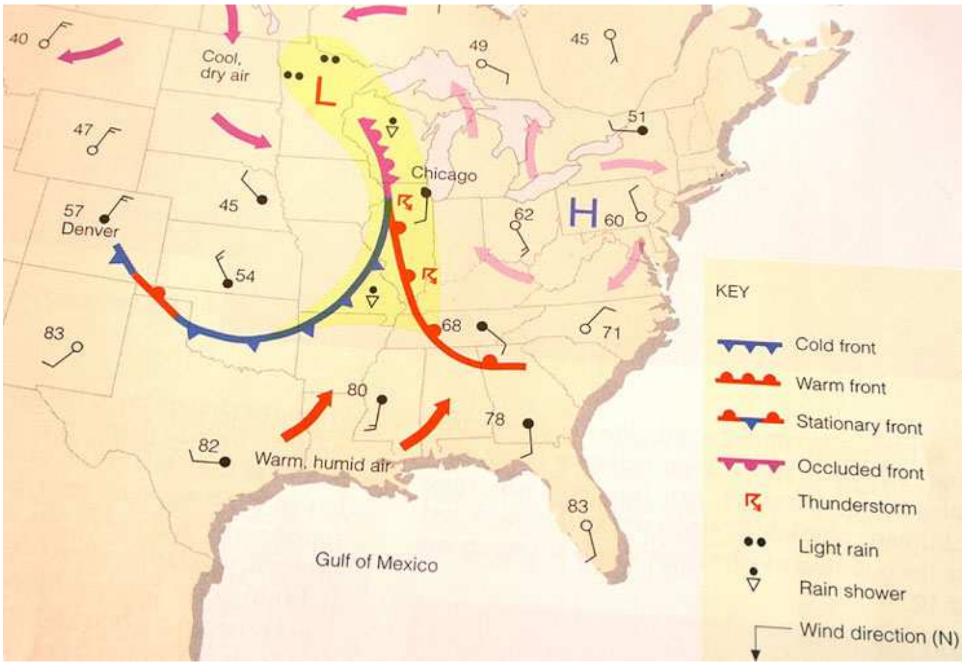

Figure 3: A weather map.

mathematical techniques taught in this course are applicable to a wide variety of practical problems.

The course will progress through several techniques. We will start with methods that use static models: line fitting, the normal equations, root finding, the Newton-Raphson method, and non-linear model fitting. We will then derive and study the Kalman filter, one of the most widely used filtering techniques for real-time filtering. The Kalman filter is only applicable to linear models and Gaussian noises. The extended Kalman filter extends the filter to handle non-linear models. The unscented transform is a recent discovery that improves the accuracy of non-linear tracking over the extended Kalman filter and can be used with some non-Gaussian noises. Finally, we will discuss the particle filter, which can be applied to any mathematical model and noise distribution.