

is the best performing technique available. Researchers in the pattern recognition community should make note of the support vector machine's many promising results if they have not already done so. This chapter will further illustrate the use of the support vector machine (SVM) with a case study addressing the use of the SVM in the context of breast cancer diagnosis from magnetic resonance imaging (MRI) examinations (subsection 4). The chapter concludes in subsection 5.

2. THE SUPPORT VECTOR MACHINE

The support vector machine (SVM) is an established supervised learning technique that creates a boundary separating two groups based on a set of provided training data [1]. Despite their strengths, support vector machines include one or more parameters that the researcher needs to fully explore in order to produce an optimal classifier. Support vector machines (SVMs) operate by locating a decision function that attempts to split the training data into two categories. The decision function is selected such that its distance to the nearest training data on either side of the surface is maximized. If no decision function is capable of linearly separating the data, a kernel transformation function is used to map the data into a different dimensional space (called a feature space) so that it can be linearly separated using standard support vector machine decision function techniques. Multiple types of kernels have been developed to map data into differing dimensions. If the kernel transformation function does not fully separate our data, a slack error variable is used to create a soft margin decision function for data separation. Figure 1 illustrates an example SVM decision function and displays the margin.




Figure 1. A support vector machine decision function (solid line) demonstrating the margin separating it from neighbouring training samples.