This thesis analyzes what acoustically sets apart recordings of healthy people from recordings of people afflicted with multiple sclerosis, and how this distinction can be used to automatically detect multiple sclerosis from fairly simple recordings of a subject's voice, potentially discovering early cases of this disease. Chapter 1 includes the theoretical background of the effect of multiple sclerosis on speech and the descriptions of the data, software, hypotheses and assumptions used here. Two sets recordings of read speech were used, a corpus of afflicted speakers and a control corpus of healthy speakers, totalling 250 individuals. A subset of this corpus was manually annotated, resulting in one dataset. Simultaneously, these entire corpora were also annotated automatically, resulting in another dataset, which was created to explore the possibility of detecting multiple sclerosis automatically. Chapter 2 describes the 13 acoustic parameters used in this thesis, their exact hypothesized relationships with the symptoms of multiple sclerosis and the ways they were calculated. Chapter 3 elaborates on the statistical testing of the aforementioned parameters, their interpretation, the success rate of the two machine learning models used to assess their total predictive power, and a potential way to apply the principles of one of these models practically. In the case of the manual dataset, all 13 parameters were used, of which 9 were statistically significant. The machine learning model trained with these data exhibited a 93% raw accuracy with a Cohen's kappa of $\kappa = 0.63$. The automatically annotated dataset was paremeterized using only 12 parameters, of which 7 were significant. The machine learning model employed on these data exhibited a raw 85% accuracy and a kappa $\kappa = 0.5$. The thesis then goes on to propose a way to practically apply these findings in a clinical setting.