The importance of quantifying the nature and intensity of emotional states at the level of populations is evident: we would like to know how individuals feel so that we may improve public policy, build more successful organizations, and more fully understand economic and social phenomena. Using the words with which we communicate, we develop an automated technique to measure the information content and happiness expressed by millions of users of the global social network Twitter. Individual messages or “tweets” originating in the United States are identified geographically in order to map the lexical diversity and happiness of Twitter users nationwide, as well as examine correlations with other spatial data. Happiness is seen to vary consistently across states, and is independent of income. Happiness and information are seen to vary with the political makeup of the US in an interesting way. Using a graphical method, we examine which individual words are responsible for these differences.