Multi-Lingual Valence Analysis Across 20th Century Literature and the Twittersphere

Eric M. Clark, Christopher M. Danforth, & Peter Sheridan Dodds

Computational Story Lab, Department of Mathematics & Statistics, Vermont Complex Systems Center, & Vermont Advanced Computing Core

Abstract

Understanding and statistically processing underlying trends in natural human language has been an ongoing goal in Computational Social Science. This work explores trends in several languages, using expressions found on the internet, in 20th century literature, and social media. We use a Hedonometer to measure happiness in several corpora, using human ratings of emotionally charged words. Previous work has established and tested the instrument on English corpora, discovering a bias towards positive word usage in billions of tweets, millions of books, music lyrics, and media articles. Until now, it has remained an open question as to whether this trend is prevalent with respect to other languages. This work extends these previous analyses through a multilingual extension of the hedonometer to uncover interesting stories and underlying trends from literature and across social media.

Measuring the happiness of a text

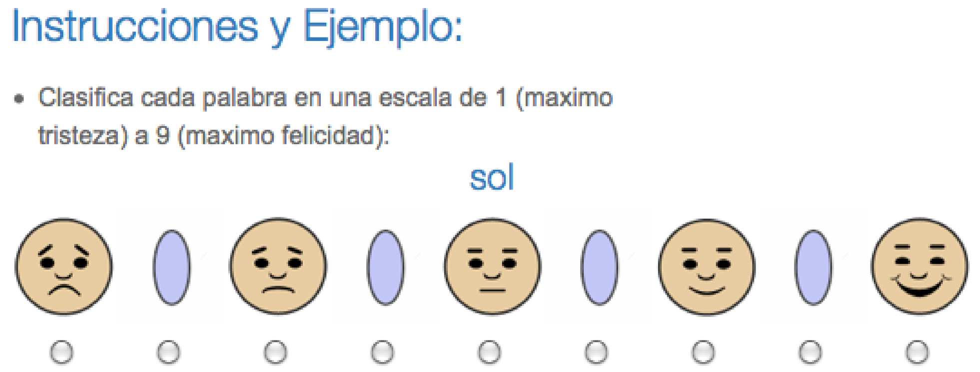

The most frequently used 10,000 words per language were compiled through frequency distributions from literature,(Google Books), websites (Google Web Crawl), and twitter. Surveys were created mimicking the self affective mannequin method, a sample of which is given above. Fifty native speakers were recruited to identify the face that best matched the emotional response elicited by each word, which were then converted to a 9 point scale. On the numeric scale, 1 corresponded to the face with the largest frown and 9 to the face with the largest smile. The average happiness score, h_{avg}, for each word is then calculated via the arithmetic mean of 50 user reported ratings per word.

Using frequency distributions from twitter and google books (right), the density of words are represented as weighted histograms for Spanish, Russian, German, French, and Chinese. The density of positive words ($h_{avg} > 5$) is significantly heavier than negative words. Since neutral articles have the highest frequency rank, the densest abundance of words appear in the neutral region, as predicted by by Zipfs law. However, each of these corpora exhibit a clear abundance of positive words over negative words which is a validation of the Pollyanna Hypothesis across language.

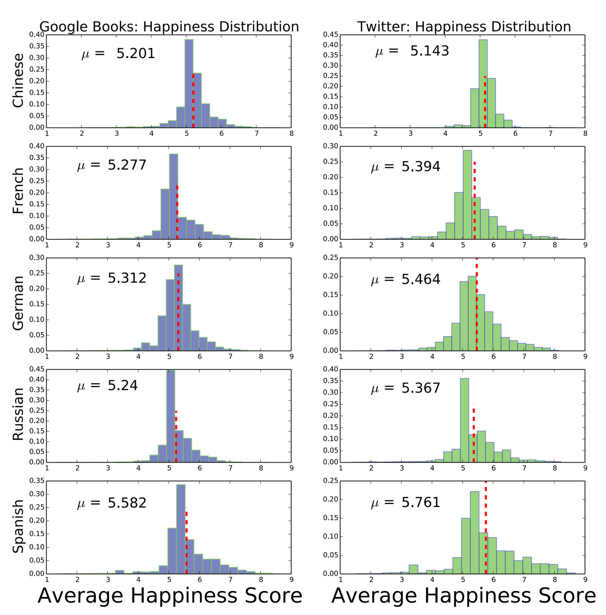

Translating Happiness

In the figure below, heat map scatterplots illustrate the correlations between evaluations of words translated across language. Each axis represents the happiness score attributed to a word by native speakers from each language. The background color represents whether the row language has an average happiness rating that is greater than the column language. Yellow represents a positive shift, and blue represents a negative shift. The number of words in common, N, the spearman correlation coefficient, R, and the average happiness shift, Δ, are displayed in the upper left corner of each subgraph. The density of points appearing on each heat map are on a color scale from white (densest regions) to black. For example, the row Spanish, column Portuguese yellow subgraph has 3273 translationally stable words, with a Spearman correlation of 0.85, and an average happiness shift of 0.07 which indicates that Spanish survey participants score words as slightly more positive on average than Portuguese participants. Each correlation is significantly positive with p-values < 0.01 , indicating that on average the emotional response of stable words is translationally invariant.

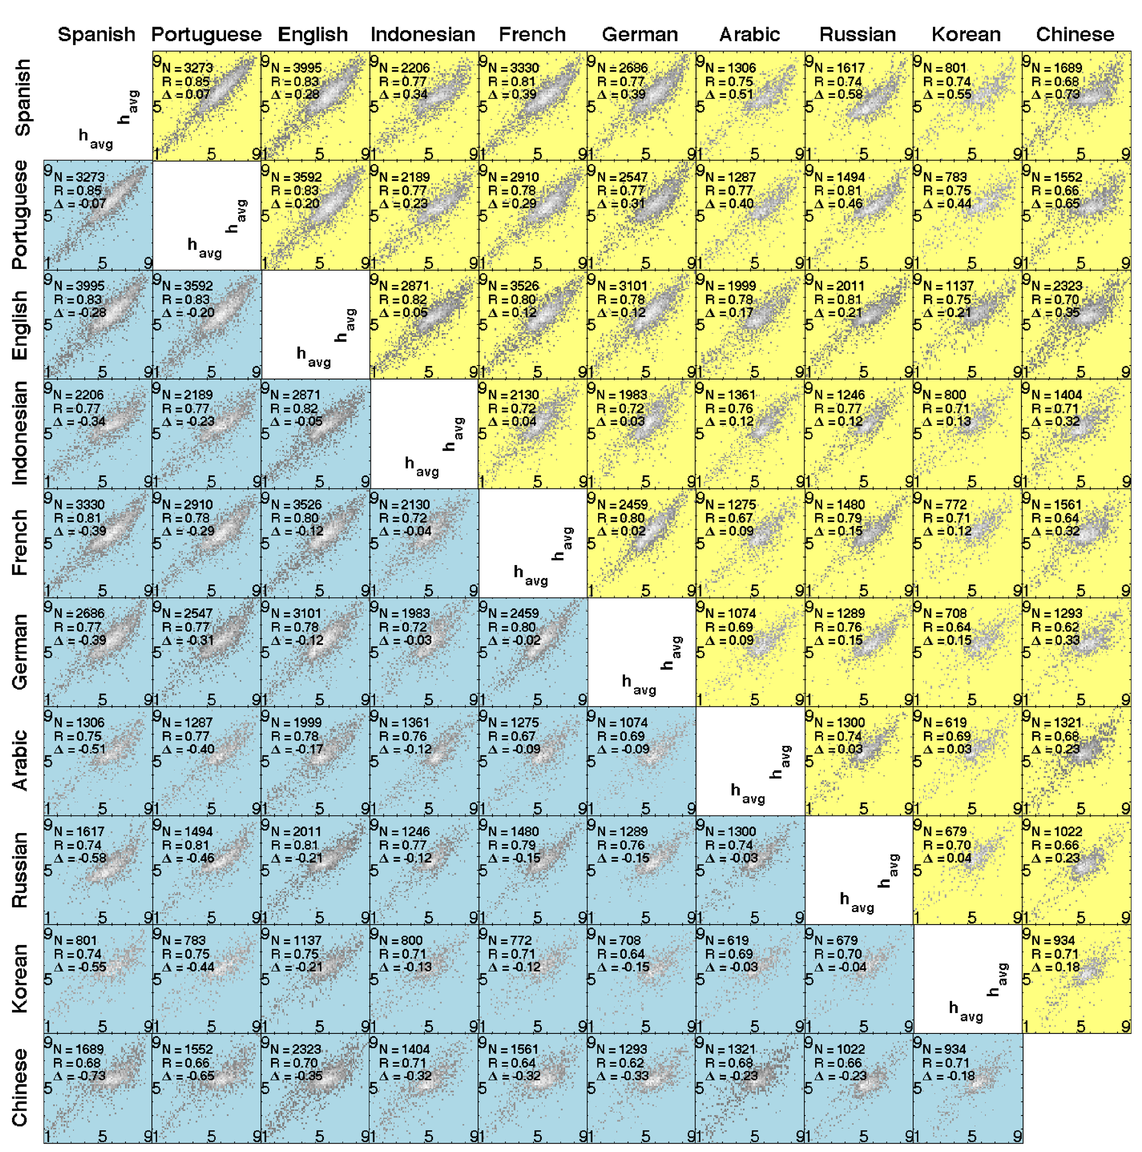

Acknowledgments

The authors wish to acknowledge the Vermont Advanced Computing Core, which is supported by NASA (NNX-08AO96G) at the University of Vermont which provided High Performance Computing resources that contributed to the research results reported within this poster. EMC was supported by the UVM Complex Systems Center, PSD was supported by NSF Career Award # 0846668. CMD and PSD were also supported by a grant from the MITRE Corporation.

Happiness Time-series of 20th Century Literature

Using the happiness scores of each wordlist distribution, the average emotional rating of a corpus is calculated by tallying the appearance of words found in the intersection of the wordlist and a given corpus. A weighted arithmetic mean of each word's frequency, f_{word}, and corresponding happiness score, h_{word} for each of the N words in a text yields the average happiness score for the corpus, $\bar{h}_{text}$:

$$\bar{h}_{text} = \frac{\sum_{w=1}^{N} f_w \cdot h_w}{\sum_{w=1}^{N} f_w}$$

The Z-scores of the average happiness value per year for 20th century literature are plotted above for English, German, French, Russian, and Spanish literature. The Z-scores of each data point help identify grave shifts from the global mean of the data set. The timeframes surrounding important historical events can be identified as outliers on the graph. Two universally negative events occur during each of the World Wars, which are indicated on the time series.

World War II Word-shifts

Word-shift graphs illustrate two separate word frequency distributions. A reference period (T_{ref}), creates a basis of the emotional words being used to compare with another period, (T_{comp}). The top 50 words responsible for a happiness shift between the two periods are displayed, along with their contribution to effecting the average happiness of the era. The reference period word distribution spans from 1900 to 1910 while the comparison word distribution spans from 1920 to 1930. The arrows (↑, ↓) next to a word indicate an increase or decrease, respectively, of the word's frequency during the comparison period with respect to the reference period. The addition and subtraction signs indicate if the word contributes positively or negatively, respectively, to the average happiness score. The words appearing on the graphs from the left are translated to English using Google's free online translation service and appear on the graphs to the right. This allows an English speaker to understand the types of words being used in the German and French corpora. Word-shift plots yield insight into the literary emotions that have been encoded in the Google Books corpus as a consequence of significant historical events.

(Right) Here we present a German word shift plot from this era. The 1920s are used as a reference period for the 1940s. Translations report an increase of 'world war', 'hitler', 'fight', 'Nazi', 'weapons' ,'adolf', 'SS', and several other combative words that are specifically related to World War II.

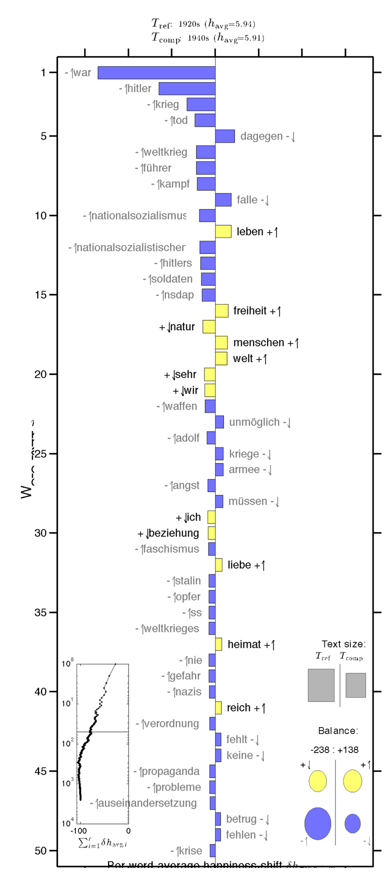

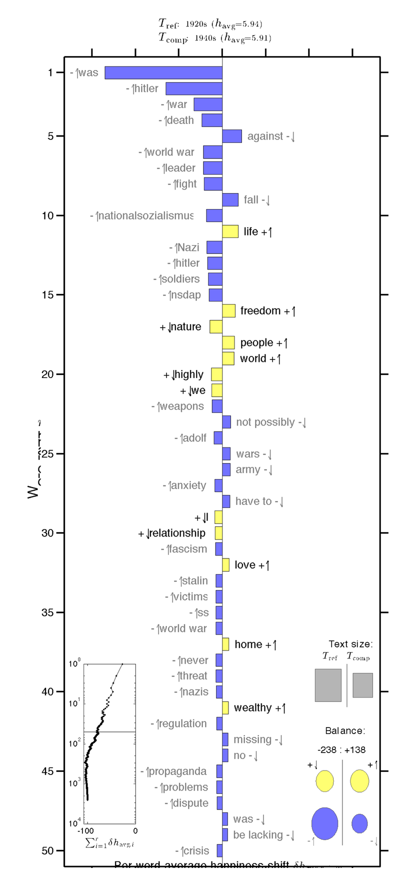

(Right) French word shifts from this era display the same theme. Here the 1920s are used in reference to compare the 1940s. Translations indicate an increase in 'hitler', 'war', 'enemy','die', 'difficulties', 'defeat' and many other indications of the war's documentation. Each of these word shifts actively convey the literary theme to match the trends in the happiness time series graphs. Understanding the emotionally charged words of each of these time-periods is an efficient way to understand the most significant contributors in swinging the emotional literary mood.

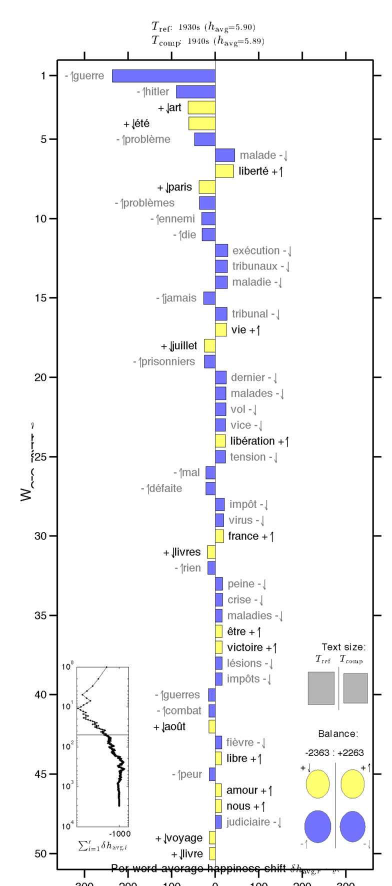

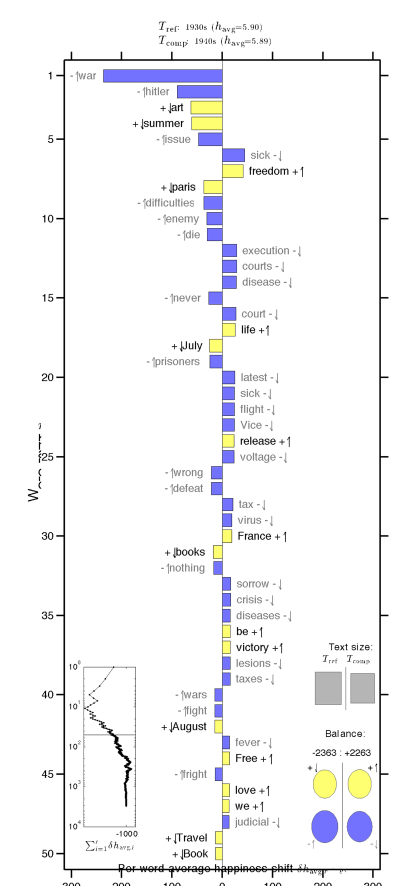

2012 Summer Olympics Happiness Timeseries

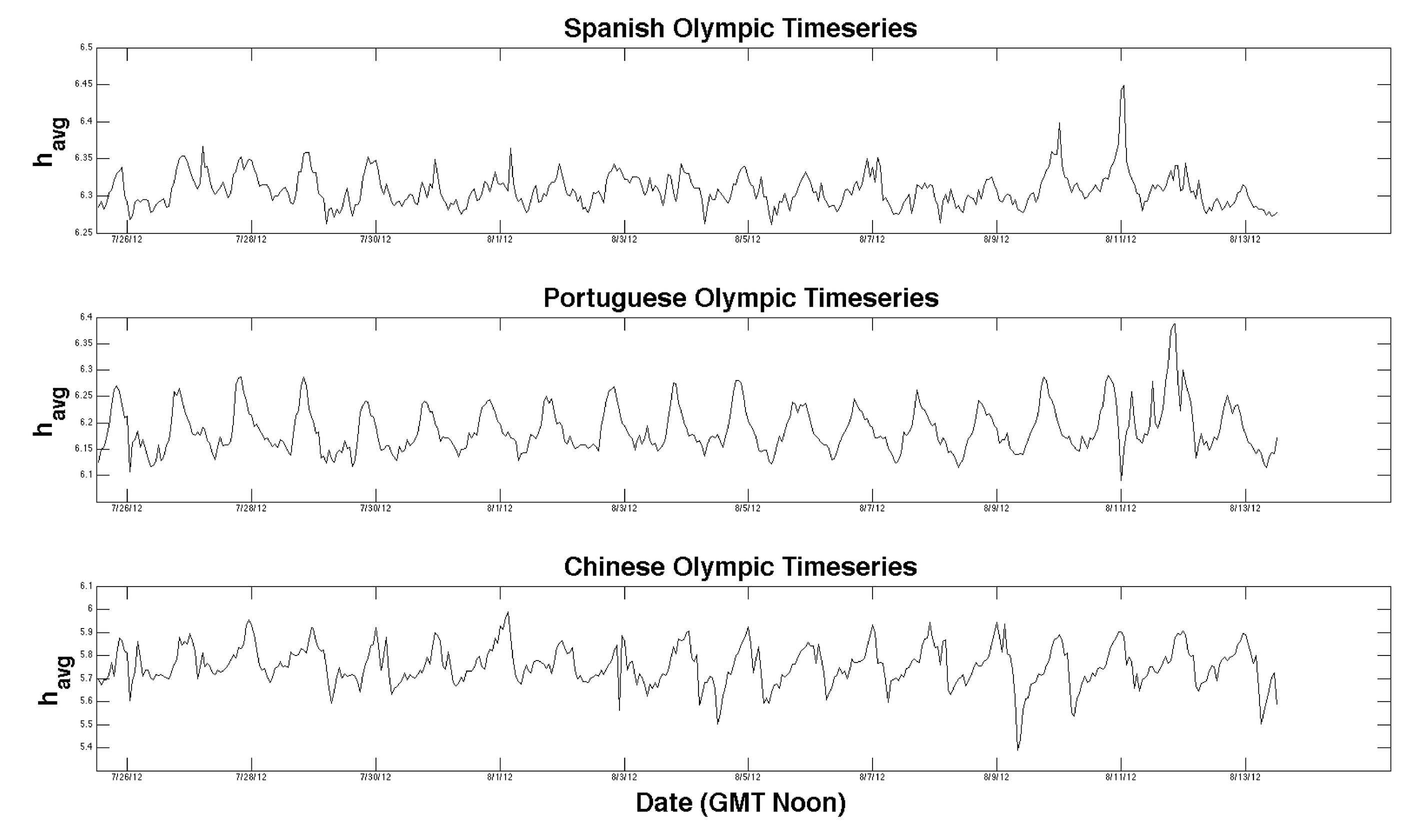

Spanish, Portuguese, and Chinese tweets are collected for each hour spanning 7/26/12 through 8/13/12. Using the hedonometer, average happiness is calculated as function of the hour of each tweet's composition. The above figure presents the happiness time series of tweets during the Olympic games.

Olympics Wordshifts

Word shift graphs convey that highly emotional tweets discussing the outcome of each olympic award are likely responsible for many of the local optima appearing on the Olympic Happiness Timeseries. In each of the following word shifts, the full Olympic data set is used as a reference period to analyze hours of interest during the Olympic games. The graphs on the right are the English translations of words appearing from the graphs on the left.

(Right) The word shift corresponding to the hour Mexico the Olympic Gold Medal in Soccer is depicted on the right. The most notable outlier on the time series plot occurs during 8/11/12 at this precise time. The words responsible for the shift include 'Mexico','gold', 'goal'. This shows the outcome of this olympic game had a dense representation twitter, and was highly responsible for the sharp incline in valence.

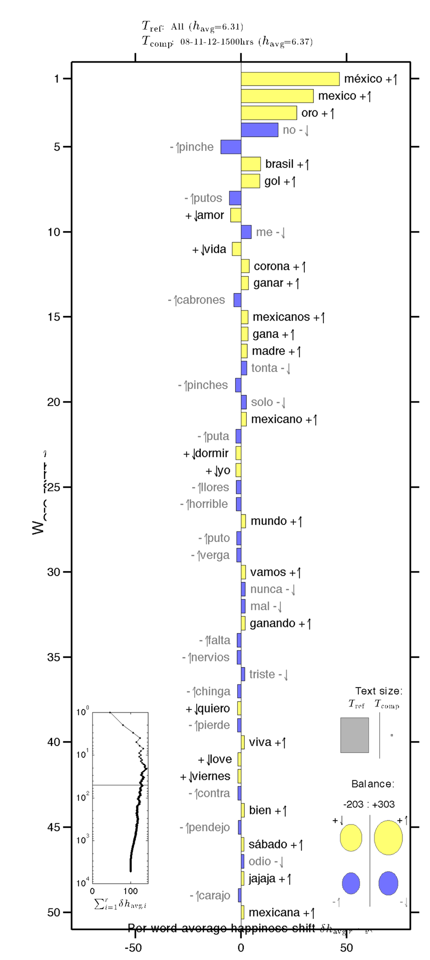

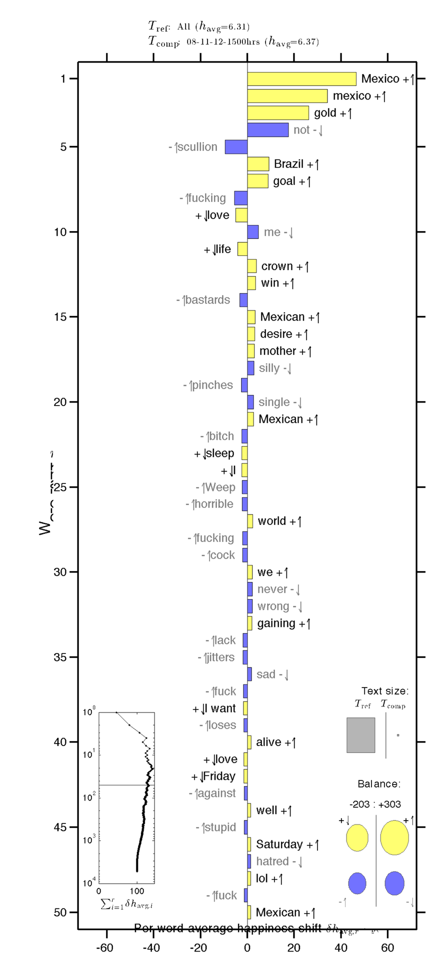

(Right) Here the word-shift of Portuguese Tweets during the 17th day of the Olympics is used for comparison against the full Olympic data set. On this day, Brazil won a gold medal in women's volleyball and a silver medal in soccer. Translations of the words appearing in Portuguese indicate an abundance of 'Brazil', 'gold', 'silver', 'goal', 'congratulations', 'volleyball', and 'football'. These words serve as an indication that the emotional shift in the time series graph is indeed due to Brazil's performance in the Olympic Games.