The Happiness Index

Research conducted and validated by third-party Data Science team, in partnership with the Prodoscore Research Council

Published May 19, 2022

prodoscore.com/prodoscore-research-council/

The Happiness Index

Why? To better understand the inner lives of employees

What is it? A metric that captures the effective ratio between the average number of positive to negative messages.

How? The model that helps generate the Employee Happiness Index was trained using over 10,000 curated chat messages from employees over a multi-year period. Document labeling was performed by the lead researcher with assistance from subject matter experts in the field of psychological wellbeing at Claremont Graduate University. Facial and Criterion validation was completed by the head researcher in collaboration with Prodoscore employees. Current categories include messages indicative of positive, negative, or neutral emotional sentiment.

Calculating the Happiness Index

Prodoscore Happiness Index Ranking	
20838	7.92
23022	6.72
20869	5.11
22295	5.06
27432	4.92
24782	4.82
25652	4.56
23023	4.55
13135	3.26
24842	3.12
980	3.07
23187	3.02
25651	0.91

Most Positive

Most Negative

In this example, the higher the number appears, the more positive the employee's general communication.

Cutoffs are determined by the standardized weekly variation in Prodoscore Happiness Scores (2.17 in this sample).

Confidential | DO NOT SHARE

Calculating the Happiness Index

Employees with Happiness Indices above their company median tend to express more positive emotions in their communication when compared to employees below the company average. Using a dichotomous/median-split allows for meaningful comparison between employee groups and provides an opportunity to assess categorical changes in employee communication sentiment.

Employees classified as most “happy” are employees who have a Happiness Index score 1 standard deviation above the company average. Employees classified as least happy are employees who have a Happiness Index below 1.0 or a score that is one standard deviation below the company average. Employees with a Happiness Index below 1.0 tend to express negative emotion at a higher frequency than positive emotion when messaging others.

Training Corpus Validation

Comparing The Employee Happiness Index Indicator to an Established NLP Library

The results below in the biserial correlation table stand as a proxy convergent and discriminant validity test of the training corpus used in the Employee Happiness Index.

Point Biserial Correlations Across the Entire Training Corpus

	LIWC Positive Emotion	LIWC Negative Emotion	Training Data Positive Emotion
LIWC Positive Emotion	-		
LIWC Negative	0.12*	- 0.15*	
Emotion	0.51*	0.37*	- 0.35*
Training Data Positive Emotion	-		
Training Data Negative Emotion	0.19*		

Training Corpus Validation

When comparing the Employee Happiness Index to the well established LIWC2015, I observed significant associations that trended in predictable patterns. Statements identified as positive by LIWC2015, demonstrated a significant positive correlation with statements labeled positive and a significant negative correlation with statements labeled negative in the training corpus. Statements identified as negative demonstrated a significant negative correlation with statements labeled positive and a significant positive correlation with statements labeled negative. The strongest relationships exists when comparing positive labels from the training corpus and LIWC2015 followed by comparing negative labels. Negative and Positive statements are not perfectly correlated (r = 1) as the training corpus contains neutral statements. Bing and NRC library results of the training dataset are available in the Appendix.

LIWC2015 Limitations on Chat Messages

The chart below highlights the classification issues with using dictionary-based approaches to NLP using LIWC2015. The following statements were assigned weight as an indicator of positive emotion by LIWC2015 despite have a negative charge.

Training Corpus Validation

Statements with 0 LIWC2015 Negative Emotional Attribution

“I met the XXXX team last week, not impressed”
“wow bunch of bs is right im with you”
“all didnt go well”
“Well this is a bust.”
“this guy is a total joke”
“that's not fair”
“my brain subconsciously does not like him”
“im not good with that”
“Not happy man!”
“I do not feel good”

When assessing natural language generated by employees on chat programs, it becomes abundantly clear that employees express themselves differently than in other mediums such as email or journaling. The word forms and morphology of sentences vary, and the semantic differences necessitate the use of a pre-trained algorithm that is hyper- tuned towards the specific task of chat message analysis.

Model Training and Performance

MODEL TRAINING

With the package collection modules such as transformers, Pytorch in Python, a BERT-base cased model was loaded and trained to classify Prodoscore employee chat message data. Data was split 80/10/10 training, validation, and testing partitions for training with a working batch size of 64.

A maximum token length of 35 was chosen based on the range of tokens present in the training partition.

Hyperparameters include:

- 5 training epochs
- Crossentropy Loss
- A weight decay implementation of the Adam optimizer (AdamW)
- An adaptive learning rate that decays by a preset amount (gamma) for every epoch

**Please contact the primary researcher for more details regarding model hyperparameters*

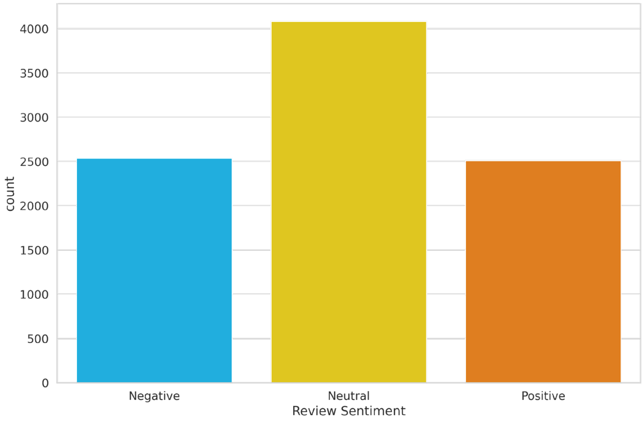

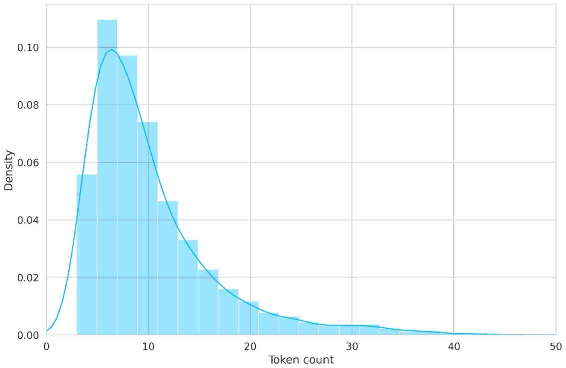

Model Training and Performance

MODEL PERFORMANCE

Testing Data

	Precision	Recall	F1-Score
Positive	0.97	0.87	0.92
Neutral	0.85	0.91	0.88
Negative	0.84	0.84	0.84
Model	-	-	88%

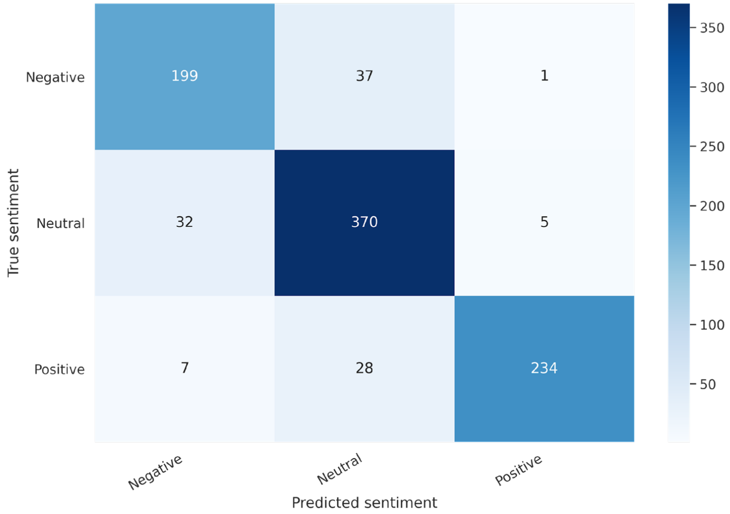

Above is a confusion matrix capturing the model performance of the Prodoscore Research Council's algorithm. The algorithm has a testing accuracy of 88% and performs relatively well at classifying chat messages across the three categories. The model excels most at classifying positive messages and rarely confuses positive messages as a negative and negative messages as positive. Future iterations will focus on optimizing the model to reduce the number of false positives (i.e., neutral/positive statements labeled negative).