Natural language processing is a field of artificial intelligence that teaches computers how to analyze and understand human language. This research is commonly seen in applications that include topics such as machine translation, chatbots, and auto-fill software. A task within natural language processing, natural language inference, is a problem in which we assess relationships between sentence pairs (neutral, entailed, or contradictory) across three diverse datasets. We enhance models with two features: negation word count and persistent homology, the latter of which is a mathematical method for spatial pattern analysis. The best negation-augmented model, RoBERTa_Neg, improved classifier accuracy by 0.18%, surpassing baselines on two of three datasets. However, persistent homology vectors had a detrimental effect, reducing accuracy by 26.50% in the best-performing model (BBU_PH_SVM) and 60.20% in the worst-performing model (BBU_PH_MLP), falling short of established baselines. Combining these features in a single classifier showed interesting results but still resulted in reduced performance compared to using the negation feature alone. All results are compared and contrasted to previously established baselines on each dataset employed in this work.