Bulletin of the *Transilvania* University of Braşov
Series III: Mathematics and Computer Science, Vol. 6(68), No. 1 - 2026, 255-272
https://doi.org/10.31926/but.mif.2026.6.68.1.17

# A MULTIDIMENSIONAL APPROACH TO PREDICTING ACTION PROBABILITY FROM TEXT[1]

**Vlad BELCIUG[2] and Elena PELICAN[*,3]**

**Abstract**

Predicting the probability of action from textual data has been an area of growing interest in various domains, including social media analysis, customer support, psychology, criminology and online marketing. Our paper proposes a novel approach that combines multiple text representation techniques, ensemble methods and a multidimensional linguistic analysis framework to predict the probability of action from text data. The approach consists of five main phases: data preprocessing, text representation, feature fusion and selection, multidimensional analysis and action probability prediction.

*2020 Mathematics Subject Classification:* 68T50, 68T30, 68T07.
*Key words:* action prediction, text analysis, multidimensional, natural language processing, synergies.

## 1 Introduction

Natural language processing (NLP) has made significant strides in understanding and extracting meaningful information from textual data. Traditional approaches, such as bag-of-words (BoW) and TF-IDF, have been complemented by advanced word embeddings and neural models ([13], [15]). However, these methods often lack the ability to capture the multidimensionality of language beyond syntactic and semantic features. This paper proposes a pipeline that combines these conventional methods with a novel dimension-based scoring system to analyze texts along distinct dimensions and the synergies between them). The aim is to provide a comprehensive framework for text analysis that can predict actions from textual data.

---

[1] *Paper presented at the 5th edition of the International Conference on Mathematics and Computer Science, MACOS-2024, June 13-15, 2024, Brasov, Romania*

[2] Faculty of Mathematics and Computer Science, *Ovidius* University of Constanţa, Romania, e-mail: vlad.belciug@365.univ-ovidius.ro

[3*] *Corresponding author*, Faculty of Mathematics and Computer Science, *Ovidius* University of Constanţa, Romania, e-mail: epelican@365.univ-ovidius.ro

## 2 Related work

Natural Language Processing (NLP) has been extensively applied in text classification, sentiment analysis and topic modeling ([2], [10]). With the development of deep learning, large language models (LLMs) such as BERT ([6]), GPT-2 and GPT-3 ([16], [3]), RoBERTa ([11], [22]), and T5 ([17]) have significantly advanced the state-of-the-art in understanding and generating human language. These models excel at capturing complex linguistic patterns and context, enabling breakthroughs in tasks like machine translation, question answering and text generation. However, despite their impressive capabilities, these models often suffer from a lack of interpretability and transparency. They operate as "black boxes, "making it challenging to understand the underlying reasoning behind their predictions ([20]). Moreover, while they capture syntactic and semantic nuances, they may not effectively model multi-dimensional features such as emotional intensity, cognitive processes, or social dynamics inherent in text ([12]). The concept of dimensional text analysis has its roots in psychological and behavioral studies, where tools like the Linguistic Inquiry and Word Count (LIWC) have been used to quantify psychological states based on word usage ([21]). These studies emphasize the importance of dimensions like affective processes, cognitive mechanisms and social relationships in understanding human communication. However, existing frameworks lack integration with advanced NLP and machine learning techniques, limiting their applicability to modern, large-scale textual data. Our work tries to bridge this gap by introducing a novel approach that integrates a multi-dimensional analysis system into a machine learning pipeline for action prediction from text. So far, action prediction efforts are limited to a few domains like next click prediction, advertising or sentiment analysis. By quantifying texts along specific dimensions—such as emotional, cognitive, temporal and social factors—and combining these with advanced text representation techniques, we hope to give action prediction a larger scope. This integration allows for the capture of latent semantic patterns through machine learning while maintaining the ability to explain predictions through interpretable features derived from the dimension system. While previous studies have addressed sentiment or topic classification, our work specifically targets the prediction of actionable texts, which has significant implications for fields like organizational communication, customer service and automated alert systems. For example, if this system proves accurate, it could be used in combination with user ratings as the foundation for creating a 'trust score' for social media users, in order to counter fake news dissimulation. By integrating the dimension system into a machine learning framework, we hope not only to improve predictive accuracy, but also enhance the model ability to provide actionable insights, filling a critical gap in current NLP research and giving a tool for further exploring the "terra nova" represented by word-action causality. The main idea is to reliably predict if an action will be triggered by a text. Since anything can be "narrated" and described using text, the larger picture seems to be that, if our model proves accurate enough, we should be able to accurately predict the probability of action reliably enough. Although our idea is at a starting point, we think that it can provide some useful insights regarding this less studied application of machine learning and AI.

## 3 Reasoning

The nature of a text is a consequence of the nature of each word than makes it up, a sum of attributes, each describing a different aspect. In order to find out the nature of the words in a text, we have used well established machine learning techniques; for example, for identifying the words that are more likely to trigger an action, we have referred to real-life causality, linguistics, psychology, marketing or customer service. As far as we know, the literature regarding this subject does not provide proven word-action causalities, but contains enough information for us to try and make educated guesses about which words would be more likely to trigger an action. For example, in psychology, trigger words are specific phrases or words that trigger emotions in the reader or listener. They are often used in marketing and advertising to influence the target audience's behavior, therefore, we can safely assume that they could be used as indicators about whether a text will trigger an action or not, along with other such "clues", we have tried to identify with this multidisciplinary approach. This is an ongoing research that constantly adds details to the description of the word-action relationship. The problem is that even if we would identify the nature of each word and we would have solid word-action causalities, we would still not be able to make a decent prediction whether the text as a whole will trigger an action or not. We need an encompassing system to infer the nature of the text as a whole, so that we can see whether the text is likely to prompt action or not. This means that we have to describe the general characteristics of the text in a quantifiable way, so that we can make decent predictions about its capability of triggering an action. For this, we used the multidimensional system - an experimental system we have created to describe textual information. In order to make it work and obtain quantifiable results, we have imagined an orchestrated pipeline to get us from text to probability of action.

## 4 Methodology

While explaining the structure of the pipeline, we briefly describe the multi-dimensional system and we detail the machine learning (ML) components.

### 4.1 Overview

In what follows, we present the pipeline of our project (see Figure 1).

**Step 1. The target data is preprocessed and the text is represented using machine learning (ML) methods**: At this step, the text data is cleaned and tokenized, stop words are removed and we perform lemmatization.

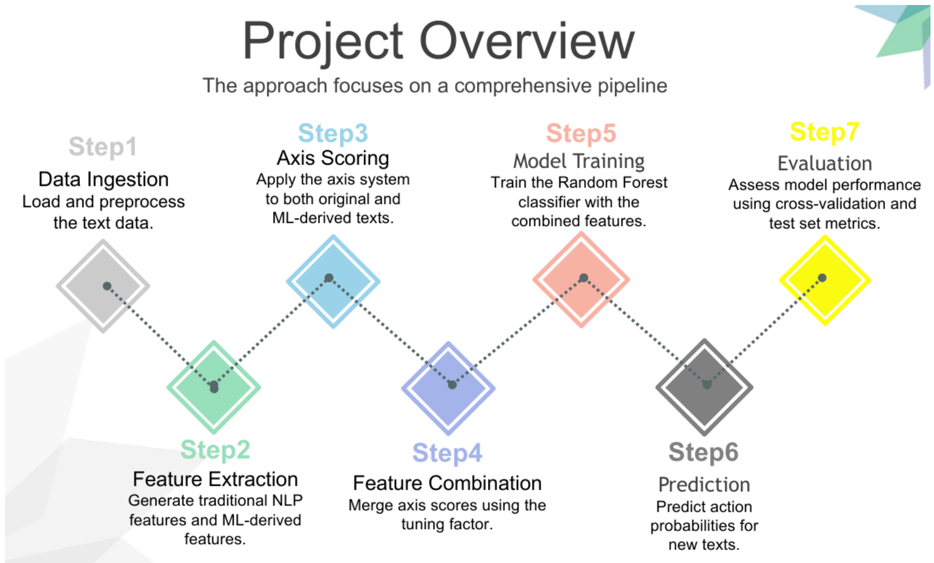


Figure 1: Pipeline Workflow

**Step 2. Initial Multidimensional System Analysis**: The multidimensional system (see subsection 4.2) is first applied to the original text to generate scores along the predefined dimensions, thus resulting the nature of the target data.
**Step 3. ML-Derived Text Features Analysis**: The target data is represented using different techniques, in an effort to explore the text from various points of view (see Figure 2).

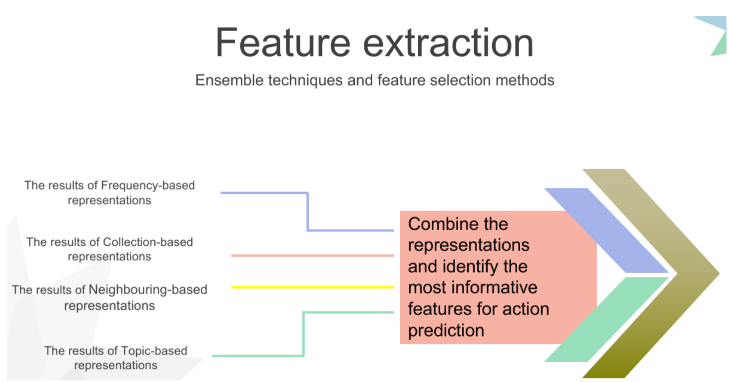


Figure 2: Feature Extraction

*Bag of Words (BoW)* is used to convert text into a matrix of token counts with CountVectorizer from the *sklearn.feature_extraction.text* module, each entry representing the frequency of a word in the document.
*TF-IDF (Term Frequency-Inverse Document Frequency)* is used to weight the frequency of words by their importance. For this we make use of the *TfidfVectorizer*

from *sklearn.feature_extraction.text*.
*N-grams* captures sequences of words (bi-grams in our case) to understand context. We implement this using *CountVectorizer* with *ngram_range=(1, 2)*. This range represents one of the future testings topics: how would more context affect the performance of the model?
*Latent Dirichlet Allocation (LDA)* is a probabilistic model that identifies topics within the text. LatentDirichletAllocation with 5 components, from *sklearn.decomposition*, is used to discover latent topics based on word co-occurrences. The number of components that LDA uses can be adjusted and so, this becomes ground for further tests.
*Non-negative Matrix Factorization (NMF)* also extracts topics, but uses a different mathematical approach. It is a deterministic algorithm which arrives at a single representation of the corpus. We use NMF command from the *sklearn.decomposition module*.
Finally, word embeddings (e.g., GloVe) are used to capture semantic relationships. We chose to use *GloVe (Global Vectors for Word Representation)* in our application because it effectively captures both global statistical information and local context of words within a corpus. This dual capability enhances our app ability to understand semantic relationships between words, which is crucial for accurate action prediction.
The transformed features from all these methods are combined into a comprehensive feature matrix (Figure 3).

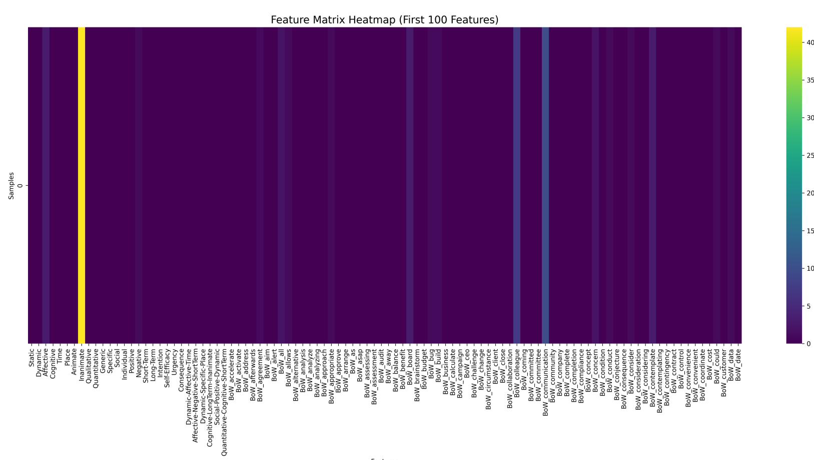


Figure 3: Feature Matrix Heatmap

This matrix serves as a high-dimensional representation of the original text, capturing various linguistic properties. The most important features (words or n-grams) from each representation are extracted. As a starting point, we have used the top 66% as a “safe” pick, in order to balance the computational needs with taking into consideration the majority of the results, while still making sure that we do not ignore important features. These top features are afterwards converted back into words to interpret the meaning of the representation and to provide input for further analysis. The top features (words) extracted from the ML representations are combined into a single text. The multidimensional system is then reapplied to this new text, producing a second set of dimensional scores. This step allows the system to evaluate how the core meanings represented by ML methods align with the original text.

**Step 4. Combining Dimension Scores**: The dimension scores from the original text and those derived from ML methods are combined. A weighted average is used, where a "tuning factor" determines the contribution of each set of scores to the final combined score (see subsection 4.2). The "tuning factor" is a weight used to control the relative influence of two sets of dimensional scores: one derived directly from the original text and another from the machine learning (ML) representations of the text (i.e., features extracted through BoW, TF-IDF, LDA, etc.). The original text dimensional scores provide a direct analysis based on linguistic markers and predefined rules, which can be very transparent and interpretable. However, they may miss deeper, more complex patterns in the text. Conversely, ML-derived features capture more nuanced relationships and interactions that might not be apparent through simple rule-based methods, but can be less interpretable. By using a tuning factor, we combine the strengths of both methods. This factor helps balance how much weight we give to the immediate textual properties versus the abstracted properties captured by the ML models. So far, we have set this to 0.5, meaning we attribute equal weight to both methods. This merging step is a crucial one, because it integrates both the raw linguistic features and the abstracted representations, providing a holistic view of the text properties.

**Step 5. Model Training and Prediction**: The combined dimensional scores form the input features for model training. Any features with zero variance (features that do not vary across data points) are removed to ensure the model only uses meaningful data. A SGD Classifier (Stochastic Gradient Descent Classifier) is used for training. This model is chosen for its ability to handle large datasets and supporting incremental learning through the *partial_fit method*. This means the model can be updated incrementally without retraining from scratch, significantly reducing the training time. The model is trained using labeled data, where each text is labeled as either "action" (likely to trigger an action) or "no action." The dimensional scores are the input features and the labels are the target. For new, unseen text, the trained model predicts the probability that the text will lead to an action. This probability is based on how similar the new text dimensional scores are to those in the training data. The rule-based probability is calculated directly from the combined dimensional scores. Each dimension score contributes to the likelihood of an action based on predefined rules or thresholds. The probability reflects the model assessment of the influence each dimensional score has on determining whether an action will be initiated. The final probability is a weighted average of the rule-based and model-based probabilities. This combination leverages both the interpretability of rule-based methods and the predictive power of the machine learning model.

**Step 6. Visualization and Reporting**: A radar chart visualizes the contribu-

tions of different dimensions and a comprehensive report is generated, detailing the findings. The radar chart, also known as a spider chart, is particularly suited for visualizing multi-dimensional data, enabling a clear and concise overview of how each dimension contributes to the overall prediction. Multiple radar charts can be used to compare different texts or datasets, showcasing how different content types or categories fare across the same set of dimensions (particularly valuable in customer service or marketing). The visual format aligns well with human intuition for understanding multi-dimensional data. It provides an at-a-glance summary that is more digestible than numerical tables or complex textual description. The pipeline is designed to be reproducible and by setting up a controlled environment (e.g. virtual environment with fixed dependencies) and using a consistent methodology, the results can be reproduced across different datasets or contexts, adding to the report's reliability. While the report is based on a fixed set of dimensions, the pipeline is flexible enough to incorporate new dimensions or ground truths as needed. This adaptability ensures that the report can be customized for different domains, enhancing its practical utility and reliability.
Together, the radar chart and the comprehensive report form a reliable tool for understanding the results of the analysis. This being said, new methods of visualization can be used and report content structures can be included to satisfy the user's particular needs. For this prototype, we found that this way of exposing the results of the analysis is satisfactory.

## 4.2 The multidimensional scoring system

The dimension-based system is a theoretical framework designed to analyze and quantify textual information across multiple dimensions. Each dimension represents a specific linguistic, psychological, or contextual property of the text. By mapping textual content onto these dimensions, the system provides a structured and comprehensive approach to understanding the multifaceted nature of language. This system is not limited to any specific application; it serves as a general tool for dissecting and interpreting text. The primary purpose of the dimension-based system is to provide measurable scores for various properties of text that are otherwise abstract or qualitative, to enable the examination of text along multiple dimensions simultaneously, capturing the interplay between different linguistic and psychological factors and offer a transparent and interpretable framework for text analysis, making it easier to understand and communicate findings.
*Brief description of the multidimensional system*: Since everything can be represented by scoring its attributes on a dimensional axis, our system uses this fundamental concept to represent a text attributes using different scored dimensions. E.g. if most of the words, as well as contexts in a text, are of a more static nature, expressing stability, passiveness, standing still, unchanging states, then the text could be safely described as static. By scoring the static dimension of a text, we could theoretically infer the probability with which the text could trigger an action, since a text containing a lot of static words suggests a lower

action probability, as it implies inactivity. The system uses different dimensions to describe text data in this manner. The dimensions are selected based on empirical research and theoretical models from fields such as linguistics, which analyzes language structures, semantics and pragmatics; psychology, which explores cognitive processes, emotions and behavior; cognitive science, which studies mental representations and information processing; social sciences which consider social context, identity and group dynamics. Studies have shown that linguistic features corresponding to these dimensions influence perception, comprehension, memory and behavior. The dimensions align with how humans process language, including attention to emotion, agency, temporality and social context. Each dimension is identified using measurable linguistic markers (e.g., word usage, syntactic structures) and can be consistently applied across different texts, allowing for comparative analysis. By assigning numerical scores to dimensions, the system enables quantitative analysis of qualitative data.

In our model, dimensions (also referred to as features) are categorized into two main groups: action-relevant dimensions and non-action-relevant dimensions, this division being critical in calculating the action probability and providing meaningful explanations based on the text analysis. The action-relevant dimensions are dimensions that, when they have high scores, *increase* the likelihood that the text will prompt an action. They are associated with language that is dynamic, emotionally charged, urgent and intentional (examples include: Dynamic, Affective,Time, Specific, Social, Positive, Short-Term, Intention, Self-Efficacy, Urgency, Consequence, synergies like "Immediate Emotional Response," "Confident Intention," etc.). The non-action-relevant dimensions are dimensions that, when they have high scores, *decrease* the likelihood that the text will prompt an action. They are associated with language that is static, cognitive, inanimate, negative, or focused on long-term considerations (examples include: Static, Cognitive, Place, Inanimate, Negative, Long-Term, Generic, Individual, Quantitative, synergies like "Cognitive-LongTerm-Inanimate, "Affective-Negative-ShortTerm," etc.).

The categorization of dimensions directly influences the calculation of the action probability and the generation of explanations in the code. During the text analysis, the code calculates a score for each dimension based on the linguistic features extracted from the text. These scores are normalized percentages between 0% and 100%. The action probability is calculated by summing the scores of action-relevant dimensions, summing the scores of non-action-relevant dimensions and calculating the action probability as a proportion (this may be subject to change as the research continues).

The dimension-based system could provide a theoretically sound and empirically supported framework for analyzing text across multiple dimensions. It offers a structured approach to quantify complex linguistic and psychological properties, facilitating deeper understanding and interpretation of textual data. Incorporating this system with machine learning methods is suitable and advantageous for tasks like predicting actions from text. The dimensions can enhance model features, improve interpretability and potentially lead to better predictive performance.

### 4.3 Machine learning and model training

The machine learning component of the pipeline is designed to enhance the predictive power of the dimension-based system by integrating it with advanced text representation techniques. The goal is to capture both explicit linguistic features quantified by the multidimensional system and the implicit patterns uncovered by machine learning models. This dual approach leverages the strengths of both interpretable, rule-based methods and data-driven, statistical learning to improve action prediction from text. After obtaining various text representations, the pipeline extracts the most significant features from each. This can be adjusted in code by selecting the top percentage of features to be considered. For our tests we have selected the top 66%. This is done by identifying the top contributing features (words or topics) for each representation (e.g., TF-IDF, LDA), based on their weights or contributions to the model. These top features are converted back into words or phrases to create a synthesized text that encapsulates the key information captured by the ML models, after which the dimension-based multidimensional system is reapplied to the ML-derived textual representations. This step aims to quantify the latent features identified by the ML models along the same dimensions as the original text, ensuring consistency in analysis. The synthesized text from the top features undergoes the same dimensional scoring procedure as the original text, generating a second set of dimensional scores that reflect the underlying patterns captured by the ML techniques. The next step is to combine the insights from both the original text and the ML-derived representations, by using a weighted averaging with tuning factor formula:

$$\begin{aligned}\text{Combined Dimensional Score} &= \text{Tuning Factor} \times \text{Original Dimensional Score}+\\ &+(1-\text{Tuning Factor}) \times \text{ML}-\text{Derived Dimensional Score}. \qquad (1)\end{aligned}$$

The tuning factor allows for adjusting the influence of each set of scores in code. A higher tuning factor emphasizes the original text features (favoring interpretability), while a lower factor gives more weight to ML-derived features (leveraging deeper semantic patterns). For our tests, we have used a 0.50 tuning factor, meaning that we attributed equal weigh for both sets of scores.

The combined dimensional scores serve as the features for machine learning. Along with individual dimensional scores, synergy scores (combinations of certain dimensions) are included to capture complex interactions that may influence action likelihood. Each synergy score counts as an additional feature, increasing the dimensionality of the feature set. Synergy scores are designed to capture complex interactions between dimensions that may not be apparent when considering individual dimensions alone. Including synergy scores can enhance the model's ability to predict actions by accounting for the combined effects of certain dimensions and by calculating synergy scores as averages of selected dimensions, we avoid creating an excessive number of features that could lead to overfitting. The feature matrix comprises all combined dimensional scores and synergy scores, providing a comprehensive representation of the text properties. A SGD Classifier (Stochastic Gradient Descent Classifier) is used for training. The dataset is split

into training and testing sets to evaluate model performance on unseen data and K-fold cross-validation is employed to ensure the model reliability and to mitigate the effects of data variance. Because actionable texts may be less frequent than non-actionable ones, leading to class imbalance, techniques such as oversampling the minority class or using class weights in the model are implemented to address imbalance. Features with low variance across samples are removed to reduce dimensionality and noise, while highly correlated features are examined to prevent redundancy. The model is evaluated based on accuracy, precision, recall and ROC-AUC to select the best hyperparameters. The trained model provides probability estimates for each class (action or no action), higher probabilities indicating a greater likelihood that the text will prompt an action. The rule-based probability is a component of the pipeline that estimates the likelihood of a text prompting an action based solely on the dimensional scores, without relying on machine learning models. This method uses predefined thresholds and heuristics to interpret the dimensional scores and compute a probability estimate and is calculated independently from the dimensional scores. Considering the fact that we score each dimension from 0 to 100, the thresholds we have set are high threshold = 80, medium threshold = 50, weak threshold = 30 meaning that scores above 80 are significant and will result in an action probability score of 100%, scores between 50 and 80 will result in 50% action probability, while between 30 and 50 will give a 50% action probability.
These thresholds represent an important fine tuning ground for the project, moving forward. By adjusting these thresholds, we can improve the accuracy of the model, based on user feedback and past performance. The Rule-Based Probability provides transparent and explainable results, as the rules and thresholds are explicitly defined and allows the incorporation of theoretical and empirical insights directly into the probability estimation. It offers an alternative estimation that can be combined with the model-based probability to enhance overall prediction accuracy.
The final probability calculation leverages both the predictive power of the machine learning model and the transparency of the rule-based system by applying the formula:

$$\text{Final Probability} = \text{Model Weigh} \times \text{Model} - \text{Based Probability}+$$
$$+\text{Rule Weight} \times \text{Rule} - \text{Based Probability}. \tag{2}$$

The code is structured with classes for text representation, dimensional scoring and model training to ensure modularity and reusability. Scikit-Learn is used for implementing classifier, grid search and evaluation metrics, while *NLTK* and *SpaCy* are used for text preprocessing, tokenization, lemmatization and POS tagging. *Gensim* is used for topic modeling with LDA and NMF.
The pipeline aims to balance the interpretability of the multidimensional system with the predictive capabilities of machine learning models. By using dimension scores as features, the model decisions can be more easily interpreted and explained, aligning with the principles of explainable AI (XAI). Machine learning

models, particularly topic models and embeddings, can capture latent semantic patterns not readily apparent through rule-based analysis. The reapplication of the multidimensional system to ML-derived features enriches the feature set with deeper linguistic insights. The design allows for the substitution or addition of different machine learning models or text representation techniques as needed. The approach can be adapted to different domains by adjusting the multidimensional system or retraining the model with domain-specific data.

# 5 Experimental setup

The dimension-based system presented in this study is an original and exploratory concept, currently in its formative stages. As a novel theoretical framework, it has not yet undergone extensive empirical validation or peer review. The system is a work in progress, intended to provide a new perspective on text analysis by quantifying linguistic and psychological properties across multiple dimensions. Recognizing its unproven status, our primary objective was to develop a functional prototype that could demonstrate the potential utility of this approach in predicting action from text. In addition to the novel dimension-based system, the combination of model-based probability (derived from machine learning algorithms) and rule-based probability (based on predefined linguistic rules) represents an innovative aspect of our methodology. This hybrid approach aims to leverage the strengths of both statistical learning and interpretability. Given its experimental nature, our focus was on establishing a functional integration of these probabilities within the predictive pipeline. We sought to assess whether this combination could enhance predictive performance and provide more nuanced insights compared to using either method alone.

A critical component of the system is the need for reliable ground truths to train and evaluate the predictive models. In the absence of universally accepted ground truth datasets tailored to our specific dimensions and synergies, we utilized a python script to randomly generate these ground truth files, with 10000 entries each of them. The script utilizes combinations of phrases to generate diverse statements made up of potentially not likely to trigger an action, the 0's and more likely to trigger an action, the 1's.

The generated ground truths serve as a provisional solution to facilitate the initial testing and validation of the system functionality.

One of the designed strengths of the system is its flexibility in allowing users to input their own ground truths that best represent their specific needs and contexts. By enabling customization, the system can be adapted to various domains and applications, enhancing its practical relevance and applicability.

To evaluate the system ability to accurately interpret affective content, we selected datasets characterized by strong emotional expressions as our target data. These datasets include texts with pronounced affective features, such as emotional narratives, opinion pieces, or user-generated content with explicit sentiments. By using affect-rich data, we aimed to test whether the system could correctly identify and

quantify the affective dimensions, as well as assess the interactions captured by the synergy scores. Affective datasets allow us to determine the effectiveness of the affective and emotional dimensions within the system, ensuring that it can detect and measure emotional content accurately. By challenging the model with data that has cleas and intense affective signals, we can evaluate the robustness of the combined model-based and rule-based probability approach in handling emotionally charged texts. The experimental setup is designed to explore the feasibility and functionality of the proposed dimension-based system and the integration of model-based and rule-based probabilities. We have focused on developing a working prototype that can be further refined and validated in future research. The use of generated ground truths and affective datasets provides a foundation for initial testing, while the system adaptability for user-defined ground truths ensures its potential applicability across different contexts.
A virtual environment is created to ensure reproducibility, with all the dependencies managed and verified before execution. System requirements, such as memory and CPU capacity, are checked to prevent resource bottlenecks.
The model performance is evaluated using standard performance metrics for a classification problem: accuracy, precision, recall and the area under ROC curve (AUC). Additionally, interpretability is assessed through feature importance analysis and the clarity of the radar chart visualizations.

## 6 Results and evaluation

We implemented the multidimensional system in a machine learning pipeline using 20000 generated ground truths (10000 for action and 10000 for no action) containing labeled texts (e.g., "action" or "no action"). We tested our approach on several emotional datasets in order to create a general impression regarding how the dimensions are being extracted, the affective nature of the texts being a valuable hint regarding some of the dimensions. The datasets include Empathetic Dialogues ([7]) (see Figure 5); Conversations Gone Awry Dataset – Wikipedia Version (CGA-WIKI)([5]) (see Figure 6); the DialogSum corpus ([4]) (see Figure 7); the "reddit_mental_health_posts dataset" ([19]) - different mental issues conversations would yield intrinsecally different results regarding action prediction (see Figure 8), as well as AI generated texts for individual dimension and synergy testing (see Figure 3).

The behavior of the model so far, as well as the results obtained using the default groundtruths, create a promising proof of concept/prototype and indicate a valid approach. The ability to customize the action prediction pipeline by incorporating user-provided ground truth data is a powerful feature that tailors the model to specific contexts and domains. This customization enhances the model predictive accuracy and relevance by aligning it with the unique characteristics of the user's data. Ground truths are labeled datasets that provide the model with examples of texts and their corresponding outcomes (e.g., whether an action

was taken or not). They serve as the foundational knowledge upon which the machine learning model learns patterns and makes predictions. By tailoring the model to the specific language, terminology and interaction patterns of a particular domain (e.g., customer service in a tech company), users increase the model predictive performance, training it on data that closely resembles the data it will encounter in practice and allowing the model to adapt to changes in communication styles, emerging trends, or company-specific protocols.
Customizing the action prediction pipeline with user-provided ground truths empowers organizations to tailor the model to their specific needs and contexts. By leveraging their own data, users can enhance predictive accuracy, gain valuable insights and improve outcomes in various applications. Whether in customer service, marketing, legal, healthcare, or internal communications, this customization feature allows the model to adapt to unique challenges and deliver actionable intelligence that drives better decision-making.

## 7 Discussion

The rapid advancement of large language models (LLMs) such as GPT-3 ([3]), GPT-4 ([14]), BERT ([6]), Gemini ([9]), Claude ([1]), Copilot ([8]), etc. has changed the landscape of natural language processing domain. These models, with billions of parameters, demonstrate impressive abilities in generating, understanding and interpreting text across a wide range of tasks. While LLMs represent a substantial leap forward in NLP, they are not without limitations. One of the main critiques of LLMs is their "black box" nature. Although models like GPT-3 can produce highly relevant and coherent outputs, they do not provide transparent, interpretable mechanisms that explain why a particular decision or prediction is made ([20]). LLMs, while powerful, are typically trained on broad, general-purpose data, meaning they may struggle in domain-specific applications without fine-tuning. In contrast, the multidimensional system can be easily adapted to specific domains by adjusting the ground truth and re-weighting the dimension scores. The state-of-the-art LLMs are computationally expensive and resource-intensive. Their implementation often requires substantial GPU power and high memory capacity, which can make them impractical for smaller organizations or real-time applications. The multidimensional system is computationally lighter, as it leverages simpler, yet effective, machine learning techniques like feature extraction and scoring. This makes it more suitable for deployment in environments where computational resources are constrained or where real-time processing is required. Hallucinations can occur when the model generates output that seems plausible but has no basis in its training data or the input prompt. This is often due to overgeneralization or confusion in the model internal representation. They can manifest as incorrect facts, invented citations, or logical contradictions. Ambiguous or unclear prompts can lead to outputs that deviate from the user's intended topic, causing hallucinations as the model tries to "guess" the right context or simply try to please the user by stating facts that are not true. The dimensional system is trainable and customizable, thus allowing constant fine tuning and perfecting. The multidimensional system addresses several key areas where current LLMs fall short. The multidimensional system breaks down text into measurable, interpretable features, metrics that are not only transparent but also meaningful to human users who can easily map the scores to actionable insights. The system proposes synergies between dimensions allowing for a more nuanced understanding of when a text is likely to prompt immediate action. All this being said, LLMs could be used in conjunction with the multidimensional system to provide deeper contextual understanding where necessary. Although the multidimensional system offers several advantages, it is not without limitations and it is in its early, experimental days. As the system relies on predefined dimensions and NLP techniques, there is always the possibility of missing subtleties in the text that an LLM could capture. Additionally, while synergies provide a more nuanced understanding of action prediction, they require careful calibration and may not generalize perfectly across all domains. As LLMs continue to improve, particularly with advancements like few-shot learning and transfer learning, their ability to handle domain-specific tasks with minimal fine-tuning will likely improve. However, the need for explainable, interpretable models, such as the multidimensional system, will remain crucial, particularly in high-stakes environments where understanding the why behind a prediction is as important as the prediction itself.

# 8 Conclusion

This paper proposes a comprehensive, multidimensional approach to predicting the probability of actions triggered by textual content. By employing a unique multidimensional system in conjunction with advanced machine learning techniques, our pipeline goes beyond traditional text analysis methods, capturing the complex interplay between language and context, hopefully providing enough information to make action become a predictable phenomenon. The initial results demonstrate that the nature of a text, whether it is dynamic or static, affective or cognitive, specific or generic, can be systematically quantified and integrated with machine learning models to predict potential actions. Our findings suggest that this method could have vast potential across various domains, from customer service or healthcare, to automated decision-making systems or even law enforcement.

Moreover, the ability to adapt the multidimensional system to specific domains, by adjusting the dimensions and their weights, highlights its flexibility. Researchers and practitioners can fine-tune the framework to suit particular applications, such as identifying persuasive language in political discourse or detecting manipulative tactics in social media content.

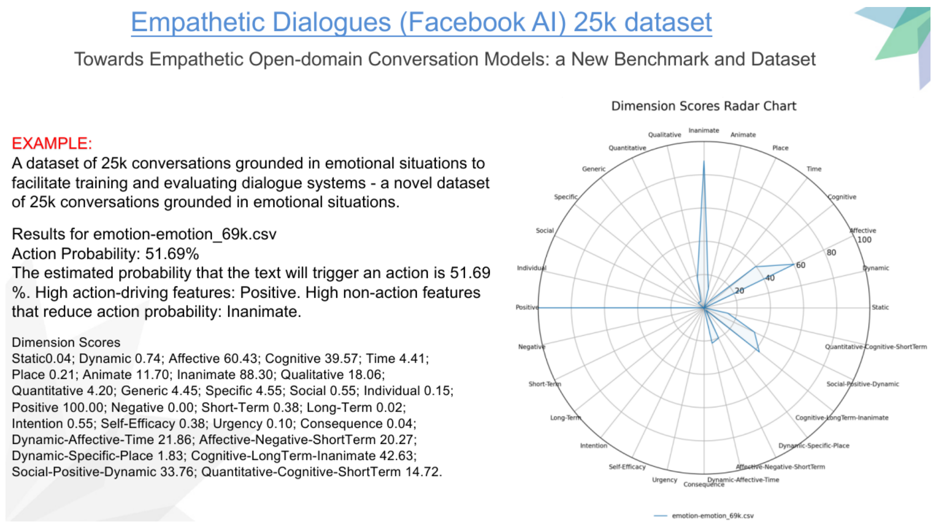


Figure 4: Dimensions scores for Empathetic Dialogues

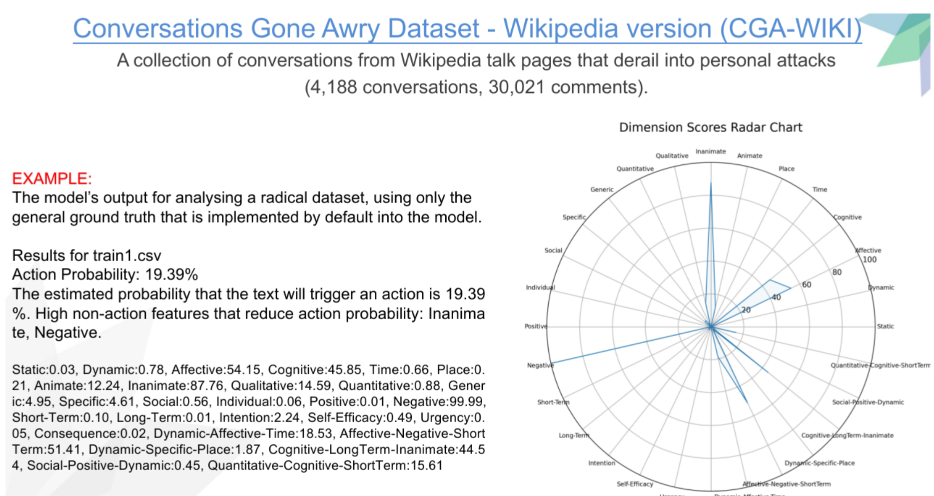


Figure 5: Dimensions scores for CGA-WIKI

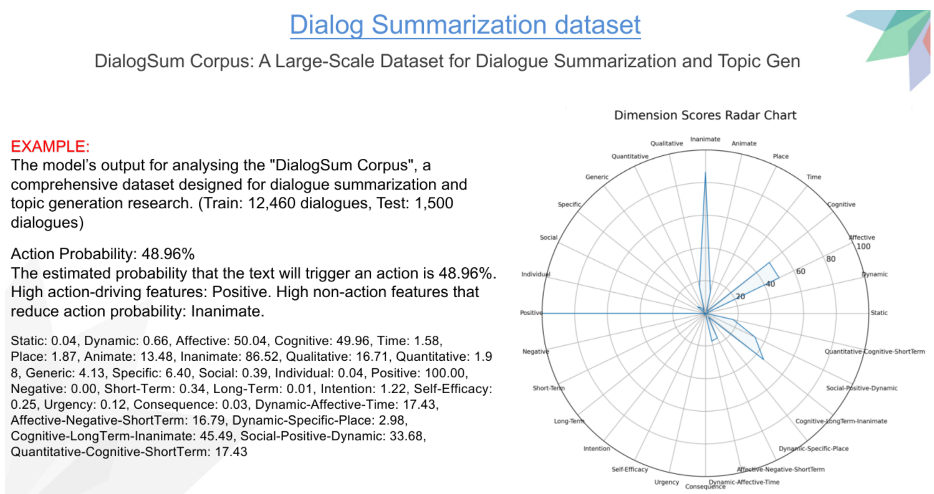


Figure 6: Dialog summarization datasets

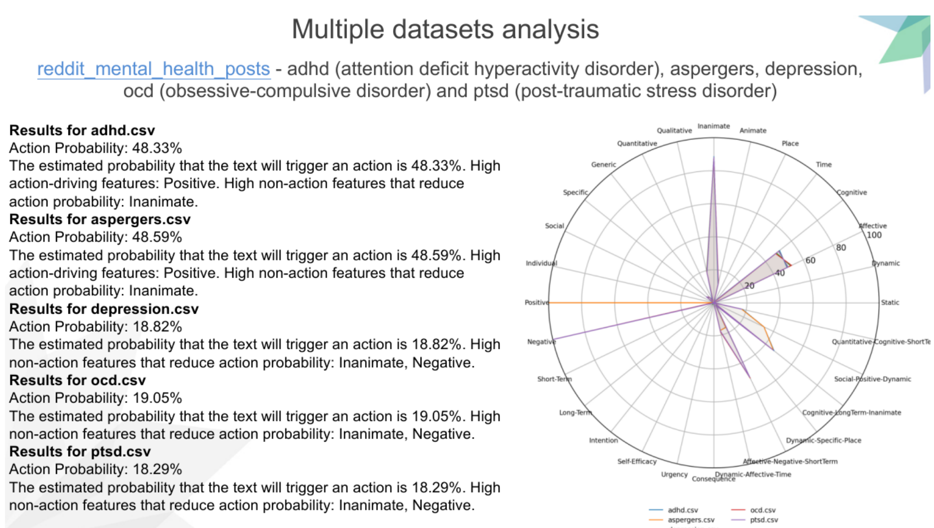


Figure 7: Multiple datasets analysis -1

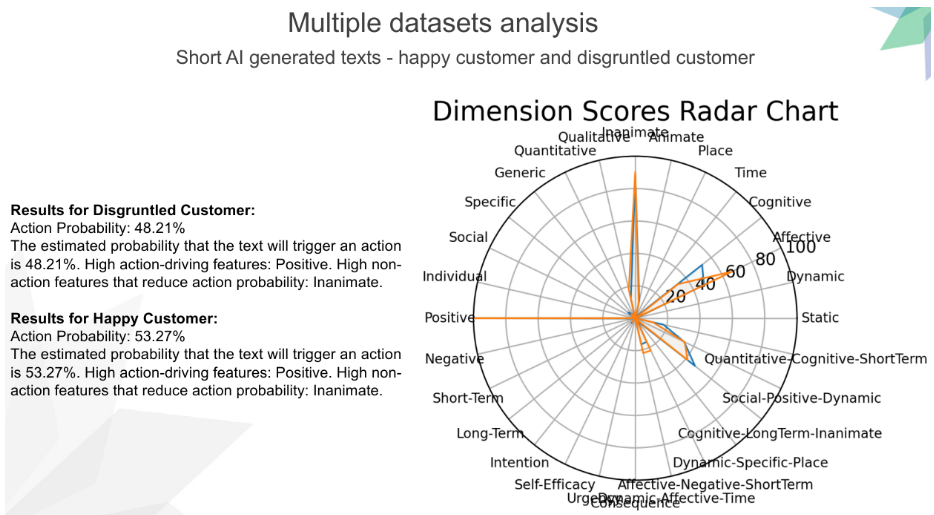


Figure 8: Multiple datasets analysis -2

## References


[1] Anthropic, *Claude AI assistant*, Software, 2023.

[2] Blei, D. M., Ng, A. Y. and Jordan, M. I., *Latent Dirichlet allocation*, J. Mach. Learn. Res. **3** (2003), 993–1022.

[3] Brown, T. B. et al., *Language models are few-shot learners*, Proc. NeurIPS (2020), 1877–1901.

[4] Chen, Y., Liu, Y., Chen, L. and Zhang, Y., *DialogSum: a real-life scenario dialogue summarization dataset*, Proc. ACL-IJCNLP (2021), 5062–5074, available at `https://www.kaggle.com/datasets/marawanxmamdouh/dialogsum`.

[5] *Conversations gone awry dataset – Wikipedia version (CGA-WIKI)*, available at `https://convokit.cornell.edu/documentation/awry.html`.

[6] Devlin, J., Chang, M.-W., Lee, K. and Toutanova, K., *BERT: pre-training of deep bidirectional transformers for language understanding*, Proc. NAACL (2019), 4171–4186.

[7] *Empathetic dialogues (Facebook AI)*, available at `https://www.kaggle.com/datasets/atharvjairath/empathetic-dialogues-facebook-a`.

[8] GitHub, *GitHub Copilot*, Software, 2021, available at `https://copilot.github.com/`.

[9] Google DeepMind, *Gemini*, Software, 2023.

[10] Liu, B., *Sentiment analysis and opinion mining*, Synth. Lect. Hum. Lang. Technol. **5** (2012), no. 1, 1–167.

[11] Liu, Y. et al., *RoBERTa: a robustly optimized BERT pretraining approach*, arXiv preprint arXiv:1907.11692 (2019).

[12] McNamara, D. S., Graesser, A., McCarthy, P. M. and Cai, Z., *Automated evaluation of text and discourse with Coh-Metrix*, Cambridge University Press, 2014.

[13] Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S. and Dean, J., *Distributed representations of words and phrases and their compositionality*, Proc. NIPS (2013), 3111–3119.

[14] OpenAI, *GPT-4 technical report*, arXiv preprint arXiv:2303.08774 (2023).

[15] Pennington, J., Socher, R. and Manning, C. D., *GloVe: global vectors for word representation*, Proc. EMNLP (2014), 1532–1543.

[16] Radford, A. et al., *Language models are unsupervised multitask learners*, OpenAI Blog, 2019.

[17] Raffel, C. et al., *Exploring the limits of transfer learning with a unified text-to-text transformer*, J. Mach. Learn. Res. **21** (2020), no. 140, 1–67.

[18] Rashkin, H., Smith, E. M., Li, M. and Boureau, Y.-L., *Towards empathetic open-domain conversation models: a new benchmark and dataset*, Proc. ACL (2019), 5370–5383.

[19] Solomonk, *Reddit mental health posts dataset*, Hugging Face, 2023, available at `https://huggingface.co/datasets/solomonk/reddit_mental_health_posts`.

[20] Rudin, C., *Stop explaining black box machine learning models for high-stakes decisions and use interpretable models instead*, Nature Mach. Intell. **1** (2019), no. 5, 206–215.

[21] Tausczik, Y. R. and Pennebaker, J. W., *The psychological meaning of words: LIWC and computerized text analysis methods*, J. Lang. Soc. Psychol. **29** (2010), no. 1, 24–54.

[22] Yang, Z. et al., *XLNet: generalized autoregressive pretraining for language understanding*, Proc. NeurIPS (2019), 5753–5763.