

2022

## Process Knowledge-infused Learning for Suicidality Assessment on Social Media

Kaushik Roy

*University of South Carolina - Columbia*

Manas Gaur

*University of South Carolina - Columbia*

Qi Zhang

*University of South Carolina - Columbia*

Amit Sheth

*University of South Carolina - Columbia*

Follow this and additional works at: [https://scholarcommons.sc.edu/aii\\_fac\\_pub](https://scholarcommons.sc.edu/aii_fac_pub)



Part of the [Computer Engineering Commons](#), and the [Electrical and Computer Engineering Commons](#)

---

### Publication Info

Preprint version 2022.

© The Authors, 2022

This Article is brought to you by the Artificial Intelligence Institute at Scholar Commons. It has been accepted for inclusion in Publications by an authorized administrator of Scholar Commons. For more information, please contact [digres@mailbox.sc.edu](mailto:digres@mailbox.sc.edu).

# Process Knowledge-infused Learning for Suicidality Assessment on Social Media

Kaushik Roy<sup>1\*</sup>, Manas Gaur<sup>1</sup>, Qi Zhang<sup>1</sup> and Amit Sheth<sup>1</sup>

<sup>1</sup>University of South Carolina, Artificial Intelligence Institute, SC, USA

{kaushikr, mgaur}@email.sc.edu, qz5@cse.sc.edu, amit@sc.edu

## Abstract

Improving the performance and natural language explanations of deep learning algorithms is a priority for adoption by humans in the real world. In several domains, such as healthcare, such technology has significant potential to reduce the burden on humans by providing quality assistance at scale. However, current methods rely on the traditional pipeline of predicting labels from data, thus completely ignoring the process and guidelines used to obtain the labels. Furthermore, post hoc explanations on the data to label prediction using explainable AI (XAI) models, while satisfactory to computer scientists, leave much to be desired to the end-users due to lacking explanations of the process in terms of human-understandable concepts. We *introduce, formalize, and develop* a novel Artificial Intelligence (AI) paradigm - Process Knowledge-infused Learning (PK-iL). PK-iL utilizes a structured process knowledge that explicitly explains the underlying prediction process that makes sense to end-users. The qualitative human evaluation confirms through an annotator agreement of 0.72, that humans understand explanations for the predictions. PK-iL also performs competitively with the state-of-the-art (SOTA) baselines.

## 1 Introduction

A long-standing problem in adopting machine learning technologies to assist humans in the real world has been the lack of a satisfactory explanation to the end-users of the technology. In the traditional machine learning pipeline, much attention is paid to fitting a function map from data points to labels. However, during the annotation of data points in the ground truth dataset, a guideline or process is often detailed by which the annotator can label the dataset. For example, to label patients for degrees of suicidal tendencies in a physical clinical setting, a well-known scale, the Columbia Suicide Severity Rating Scale (CSSRS) [Bjureberg *et al.*, 2021], is used to determine the right set of labels. Figure 1 shows

this scale. Thus, it is clear to the patient how a particular suicidal tendency is recognized once the clinician evaluates the questions and patient responses.

Similarly, when data points in a dataset are annotated in other domains, each data point is evaluated against a process or guideline similar to the CSSRS by several human annotators. The assumption is that the machine learning algorithm will implicitly recover the underlying process or guideline used by the annotators when learning a function map from data point to label. Popular methods for XAI such as LIME and SHaP, are used to explain the function learned, often through local approximations related to a single or sampled set of data points [Adadi and Berrada, 2018; Ribeiro *et al.*, 2016; Lundberg and Lee, 2017]. However, due to the black-box nature of the function and the non-convexity of the hypothesis function surfaces, it is challenging to evaluate if the recovery of the underlying process or guideline was successful and is meaningful to the end-users. Fundamentally, we might think of these XAI methods as trying to roughly construct an explanation as saying, “This data point is explainable using a simpler hypothesis function (a local approximation) due to *similar* data points (data points in the local neighborhood) also being classified correctly by the simpler hypothesis”. Consequently, much depends on the choice of local approximation and the machine learning models understanding of *similar* data points, on what is already a highly non-convex gargantuan function such as a large SOTA language model (LM) [Vaswani *et al.*, 2017]. Also, while such an explanation may satisfy the computer science community, “similarities” are hardly adequate for the end-user (e.g., a psychotherapist). The pertinent questions include: Would the human annotators consider the data points deemed *similar* by the LM also to be similar to each other? Would the human annotators agree that the explanation by the local approximation is aligned with the process or guideline used by them to label those data points?

In our study, we ask the question, what if we were able to use not just the annotator’s labels, but also the process or guidelines used to label them and explicitly control the learning of a model to recover the process or guideline (instead of implicitly). Such an algorithm would, by design, be explainable and emulate the humans model of *similarity* between data points. This paper takes the first step to answer this question grounded in the deep learning task of suicidal-

---

\*Contact Author

<b>SUICIDAL IDEATION</b>		Since Last Visit
Ask questions 1 and 2. If both are negative, proceed to "Suicidal Behavior" section. If the answer to question 2 is "yes", ask questions 3, 4 and 5. If the answer to question 1 and/or 2 is "yes", complete "Intensity of Ideation" section below.		
<b>1. Wish to be Dead</b> Subject endorses thoughts about a wish to be dead or not alive anymore, or wish to fall asleep and not wake up. <i>Have you thought about being dead or what it would be like to be dead?</i> <i>Have you wished you were dead or wished you could go to sleep and never wake up?</i> <i>Do you wish you weren't alive anymore?</i> If yes, describe:	Yes No <input type="checkbox"/> <input type="checkbox"/>	
<b>2. Non-Specific Active Suicidal Thoughts</b> General, non-specific thoughts of wanting to end one's life/commit suicide (e.g., "I've thought about killing myself") without thoughts of ways to kill oneself/associated methods, intent, or plan during the assessment period. <i>Have you thought about doing something to make yourself not alive anymore?</i> <i>Have you had any thoughts about killing yourself?</i> If yes, describe:	Yes No <input type="checkbox"/> <input type="checkbox"/>	
<b>3. Active Suicidal Ideation with Any Methods (Not Plan) without Intent to Act</b> Subject endorses thoughts of suicide and has thought of at least one method during the assessment period. This is different than a specific plan with time, place or method details worked out (e.g., thought of method to kill self but not a specific plan). Includes person who would say, "I thought about taking an overdose but I never made a specific plan as to when, where or how I would actually do it...and I would never go through with it." <i>Have you thought about how you would do that or how you would make yourself not alive anymore (kill yourself)? What did you think about?</i> If yes, describe:	Yes No <input type="checkbox"/> <input type="checkbox"/>	

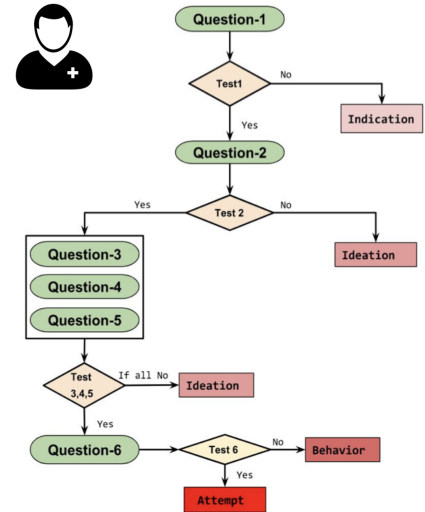


Figure 1: *Left*: Columbia Suicide Severity Rating Scale. *Right*: The induced process knowledge.

ity assessment from social media data, where incorporating the knowledge of medical processes and guidelines is of critical considerations. To this end, we propose a novel class of algorithms **Process Knowledge infused Learning (PK-iL)** for suicidality assessment from social media. We make the following contributions:

- Define **Process Knowledge (PK)** and create a dataset for suicidality assessment task based on CSSRS with annotations to include both PK and labels.
- Develop **Process Knowledge infused Learning (PK-iL)**, an explainable algorithm that explicitly controls the learning model to recover the process by effective utilization of PK in the annotation and a globally optimal optimization objective.

**Can PK-iL utilize SOTA LMs?** We note that the notion of *similar* data points is the machine learning model’s way of understanding the human annotator’s annotation process, i.e., fundamentally, the goal of the model is constructing a similarity space that mimics the human annotator’s understanding. In many domains and applications, large and SOTA LMs have excelled at capturing the similarity of some examples exceptionally well. Hence, we believe that rather than try to implicitly learn a *similarity* space as a model of the human annotator’s understanding over the whole space of examples, we can leverage SOTA LMs to define the annotator’s understanding at process-specific checkpoints in the **PK**. For instance, if a **PK** has five questions (or guidelines) to go through, we can use the SOTA model as a proxy to understand if a human annotator would have judged the guideline as satisfied. Such finer-grained understanding can potentially leverage the ability of SOTA models to learn similarity spaces while still maintaining the explicitly explainable **PK-iL** structure.

We will see the data collection, examples and intuitions, and the formalization of **PK-iL** in action through the following sections.

## 2 PK enhanced CSSRS dataset

Before the pandemic, suicidality was already a leading mental health issue across the world. Since the pandemic, incidents of suicidality have increased even further. Thus for both demonstrating high real-world impact through an important use case on real data and users as well as for ease of exposition, we explain our methods and experiments anchored around the application of suicidal thought pattern detection. However, **PK-iL** is generalizable to any domain that requires the integration of **PK** with data to derive high-quality explanations.

To conduct our study in a physical, real-world experimentation setting, we require responses from users physically present during the experiment. Consider the clinical setting of suicidal thought assessments using the CSSRS - obtaining access to a physical clinical setting presents many hurdles such as ethics approval, incentives for honest responses, etc. The demand-supply deficit in mental health already makes it hard to find a quality experimental setting, and the recent COVID-19 pandemic has compounded this issue. The significant number of persons turning to social media platforms presents an exciting opportunity to leverage a large amount of data as a proxy for user responses. Thus we utilize the dataset of Gaur et al., [Gaur et al., 2021; Alambo et al., 2019] which uses the CSSRS to label user posts from suicide-related subreddits and thus provides a real-world test-bed to evaluate the performance and explainability of **PK-iL**. Through the CSSRS, the domain experts in the study annotated longitudinal data from 448 users for the following labels: Suicide Ideation, Suicide Behavior, Suicide Attempt. High standards in annotation were maintained with a substantial inter-rater agreement of 0.84. Crucially, we expand this dataset to include the specific guideline (**PK**) for annotation in addition to the label. Table 1 shows examples of the dataset expanded with **PK**.

User text (x)	Process Knowledge (PK)	CSSRS Label
[...] a voice telling me to kill myself [...] yes - I think I should do it [...]	1.2 (yes - a voice telling me to kill myself) 2.2 (yes - I think I should do it) 4 (yes - I think I should do it)	Behavior or Attempt
[...] Rarely is a day where I dont suffer from thoughts of self-harm...	1.1 (yes - Rarely is a day where I dont suffer from thoughts of self-harm) 2 (no - no words indicating active suicidal thought)	Ideation

Table 1: Examples of data set annotation expanded with **PK**. The [...] collapses the rest of the post for brevity. Each question (1-6) in the CSSRS has a main question and sub questions 1.1, 1.2 etc, as can be seen in Figure 1. Thus the **PK** denotes the main question or sub-question being answered in the user’s Reddit post.

### 3 Tree Structure of PK

We see in Figure 1 that **PK** can be viewed as a decision tree. A process tree (**Process Knowledge**(PK)) to determine the probability of a label  $y$  for a user post can be written as a polynomial of the form

$$y = \sum_{l \in \text{Leaves}} p_l \prod_{i=1}^{N_q} (I_{yes}(q_i))(1 - I_{no}(q_i)) \quad (1)$$

, where  $N_q$  is the number of questions in the decision tree,  $I_{yes}(q_i)$  and  $I_{no}(q_i)$  represents if the post follows a *yes* path or a *no* path to the question  $q_i$ . *Leaves* is a set of all leaves that lead to the label  $y$ . For example, there are two paths in Figure 1 that lead to  $y = \text{Ideation}$ . Here  $p_l$  is computed as the ratio of the number of annotators that chose that path for the example to the total number of annotators - this in some sense captures the inter annotator agreement for those set of examples. For example, considering a particular post, if among three annotators, two annotators labeled the **PK** as the path  $1.2 \rightarrow 2.2 \rightarrow 4$ . Then the probability of  $y = \text{Behavior or Attempt}$  for that post is 0.66. Note here that the sub questions aren’t stored in the tree leaves. The path  $1.2 \rightarrow 2.2 \rightarrow 4$  is equivalent to  $1 \rightarrow 2 \rightarrow 4$ . This is done for all the examples in the training set and the final probability is an average of all the examples.

**Assertion.** For any model  $\mathcal{M}(y)$  that approximates the probability of label  $y$  for a post according to Equation 1, let the inter-annotator for the post labeled as  $y$  be  $\mathcal{A}(y)$ . Then best approximation for the post,  $\mathcal{M}^*(y) \leq \mathcal{A}(y)$

We claim the above as an assertion instead of a theorem as it is trivial to see that  $\prod_{i=1}^{N_q} (I_{yes}(q_i))(1 - I_{no}(q_i)) \leq 1$  always, and therefore any approximation is upper bounded by the inter-annotator agreement. This makes intuitive sense as what we are really interested in capturing the annotator’s thought process while labeling and on par accuracy. Improving upon the inter-annotator agreement may mean capturing something that is not present in the ground truth. Thus, we are interested in labeling unseen data as well as the human annotators would while explicitly capturing their annotation process in the learned model.

How do we define mathematically that a post follows the *yes* path or the *no* path to a question  $q_i$ , i.e., we need to define exactly what  $I_{yes}(q_i)$  and  $I_{no}(q_i)$  means. Recalling our un-

derstanding of similarity between question and answer as being a proxy to answered as *yes* or *no*, we can use inner product based similarity between representations of question and post to determine  $I_{yes}(q_i)$  and  $I_{no}(q_i)$ . For example, for a similarity model that takes as inputs representations for “*Have you thought about being dead or what it is like to be dead*” and “*Rarely is a day where I don’t suffer from thoughts of self-harm*”, the output is a value indicative of high similarity relative to other input pairs. This is seen as the question “*Have you thought about being dead or what it is like to be dead*” being answered as *yes* by the response “*Rarely is a day where I don’t suffer from thoughts of self-harm*”. There are several options in Natural Language Processing (NLP) literature to construct representations of text

- **Count Vectorizer.** Each sentence or text fragment is represented as counts of the words in the fragment padded with zeros according to largest fragment. Count vectorization however, does not consider the importance of words across different parts of the post, for example stop words might occur most frequently but provide little context.
- **TF-IDF.** TF-IDF corrects the deficiency of the Count Vectorizer method by adjusting counts by weighting for contextual importance across the post. However, TF-IDF still relies on exact matches of words being present or absent in the post.
- **Hashing Vectorizer.** Each sentence or fragment of the text is simply passed through a hash function. The idea is that similar fragments produce similar hash codes. The cryptic nature of the hash function (this is by design) is not amenable for interpretation or explainability analysis of the learned function.
- **Text Embeddings.** These are a set of neural network models that represent text in a vector space. Models such as word2vec, Transformer LMs such as GPT-3 and BERT are all examples of large neural networks that map the text to a vector space such that contextually similar texts are placed close together in the vector space while dissimilar texts are placed apart [?; ?; ?]. Since, these are the state of the art and have shown remarkable efficiency and performance in recent years, we will use Transformer based LM representations for the text.

Note that the Text Embedding models provide vector representations of words. To construct a representation for the text fragment one might average the word representations contained in the fragment. However, this loses information about the order of the words and phrases in the text and hence we use a concatenation representation padded with zeros according to the longest text fragment. Thus generally we will denote a similarity function by  $\mathcal{K}$  and representations of text  $x$  and question  $q_i$  using an embedding model as  $x^R$  and  $q_i^R$  respectively. Thus

$$\left( \pm \mathcal{K} \left( \frac{x_{sub}^R}{|x_{sub}^R|}, \frac{q_i^R}{|q_i^R|} \right) \geq \pm \theta_i \right)$$

denotes the similarity between the text and question where  $\theta_i$  are suitably chosen thresholds of accepted high similarity. The normalization of the representations by size is what makes an inner product a valid similarity measure in the range  $-1$  to  $+1$ . We will now formally develop the algorithm for **PK-iL**

## 4 The PK-iL Algorithm

We define a function that predicts the probability of post label being  $Y = y$  according to the **PK** as follows:

$$P(Y = y | X = x) = \sum_{l \in Leaves} p_l \prod_{i=1}^{N_q} \mathbb{1}_{x_{sub} \in x} \left( \pm \mathcal{K} \left( \frac{x_{sub}^R}{|x_{sub}^R|}, \frac{q_i^R}{|q_i^R|} \right) \geq \pm \theta_i \right) \quad (2)$$

, where  $p_l$  is defined as detailed in Section 3,  $x_{sub} \in x$  is a fragment of the post  $x$  (For example a sentence).  $x_{sub}^R$ , and  $q_i^R$  are representations of  $x_{sub}$  and  $q_i$  from an embedding model, and  $\mathcal{K}$  is a inner product function to measure similarity.

$\pm$  signifies if we are checking if the question  $q_i$  is answered as *yes* or *no* by fragment  $x_{sub}$  in post  $x$  with confidence  $\theta_i$ . Using  $\bigvee_{k=1}^K z_k = (\sum_{k=1}^K z_k \geq 0.5)$ , we have:

$$P(Y = y | X = x) = \sum_{l \in Leaves} p_l \prod_{i=1}^{N_q} \sum_{x_{sub} \in x} \left( \pm \mathcal{K} \left( \frac{x_{sub}^R}{|x_{sub}^R|}, \frac{q_i^R}{|q_i^R|} \right) \geq \pm \theta_i \right) \geq 0.5 \quad (3)$$

We can then optimize the Bernoulli Loss  $\mathcal{L}$  for an input post  $X = x$  and label  $Y = y$  is as follows:

$$\begin{aligned} \mathcal{L}(\{\theta_i\}_{i=1}^{N_q}) &= P(Y = y | X = x) \log(P(Y = y | X = x)) \\ &+ \\ &(1 - P(Y = y | X = x)) \log(1 - P(Y = y | X = x)) \end{aligned}$$

We perform hyperparameter tuning to choose the embedding model, fragment size  $x_{sub}$ , and  $\mathcal{K}$  (see Section 5). Since  $\mathcal{L}(\{\theta_i\}_{i=1}^{N_q})$  is strongly convex, we use Newton’s optimization method to learn the parameters of the model. The algorithm for **Process Knowledge infused Learning (PK-iL)** is as follows: Here we see that **PK-iL** is general enough to allow embedding models suitable to the task and **PK** suitable to any domain. However, in our experimental results we will evaluate **PK-iL** both quantitatively and qualitatively evaluation using the expanded **PK** enhanced CSSRS dataset (see Section 2).

---

### Algorithm 1 Process Knowledge infused Learning (**PK-iL**)

---

- 1: Compute  $p_l \forall$  leaves  $l$  from the ground truth
  - 2: Choose Kernel  $\mathcal{K}$ , fragment size, and **CE** model for representation
  - 3: Initialize  $\theta_i, \forall i \leftarrow 1$  to  $N_q$
  - 4: **for**  $k \leftarrow 1$  to  $K$  **do** ▷ Begin Newton’s method
  - 5:     **for**  $\theta_i$ , where  $i \leftarrow 1$  to  $N_q$  **do**
  - 6:         Compute  $\theta'_i = \nabla_{\theta_i} \mathcal{L}(\theta_i)$
  - 7:         Compute  $\theta''_i = \nabla^2_{\theta_i} \mathcal{L}(\theta_i) = \nabla_{\theta_i}(\nabla_{\theta_i} \mathcal{L}(\theta_i))$
  - 8:         Set  $\theta_i = \theta_i - \frac{\theta'_i}{\theta''_i + 1}$  ▷ add 1 to avoid divide by zero error
  - 9: **return**  $\theta_i, \forall i \leftarrow 1$  to  $N_q$
- 

**Prediction:** Prediction is carried out by choosing the summand in Equation 3 that has the highest value once normalized by dividing by the sum of the summands, in order for it to be a probability.

## 5 Experimental Setup and Results

For the LM to understand language in the context of suicidal thought patterns it needs to be fine-tuned on such a data. For this we word2vec representations on corpus of suicide related subreddits as well as fine-tune LMs during training on the same corpus. Thus we obtain embeddings of the text contextualized to suicidal conversation in order to accurately infer *yes* or *no* from similarity. To implement the word2vec model, we use the gensim library and the Continuous Bag of Words (CBOW) model [Mikolov *et al.*, 2013]. Note that in the word2vec model due to lack of tokenization coverage as in LMs, we chunk the string one letter at a time and check against the list of words and their vectors. The LMs we fine-tune are:

- **XLNET** - An auto-regressive language model in which the training objective calculates the probability of a token conditioned on all permutations of tokens in a fragment. When trained on a very large data, the model achieves SOTA performance across several tasks in the GLUE benchmark [Yang *et al.*, 2019][Wang *et al.*, 2018]. We use the default hyperparameters during training.
- **RoBERTa** - A BERT based model where the hyperparameters are further tuned for superior performance over BERT on GLUE benchmark tasks. We use the default hyperparameters during training RoBERTa [Liu *et al.*, 2019].
- **Google T5** - A transformer model trained on various text-to-text tasks such as translation, summarization, etc. T5 demonstrate superior transfer capabilities across tasks in the GLUE benchmark. We use the T5-small model and the default parameters [Raffel *et al.*, 2019].
- **ERNIE** - A transformer model trained on natural language corpora and large knowledge graphs and is thus a suitable model to consider in a structured knowledge intensive field like mental health. We use the default parameters for ERNIE [Zhang *et al.*, 2019].

- Longformer - A transformer model that excels at capturing long text inputs. As some of the posts can be over 8000 characters long, the longformer is a suitable model to consider for our dataset. We use the default parameters for the long former [Beltagy *et al.*, 2020].

We believe the wide range of transformer architectures above are sufficient to test our approach. We train all our models on the Google Colab platform.

**Inner Products** Bubeck *et al.*, shown that  $O(nd)$ , where  $n$  is the number of data points and  $d$  is the true underlying data dimension [Bubeck and Sellke, 2021]. The Transformer outputs are already high dimensional but Bubeck show that for natural language the models still need to get larger! Thus we use a popular trick to compute inner products in higher dimensions - the Kernel trick. Polynomial Kernel can project the data to very high dimensions and the Gaussian Kernel can project the data to an infinite number of dimensions. We see the use of a Kernel significantly improves the performance over simple cosine similarity (polynomial kernel of dimension 1). In our experiments we use the Gaussian Kernel to compute the inner product.

For the fragment size we found a span of 1-2 sentences to be the best performing model for each transformer and kernel choice.

## 5.1 Quantitative Evaluation

For baseline accuracy we directly use the embedding models to predict the label as in a traditional machine learning pipeline. For word2vec, we use the representations of the post and pass it through a logistic regression model. We make a slight modification where weights for all entries for a single word vector are shared. Table 2 shows a comparison of accuracy for all the models with their baseline, **PK-iL** with Cosine Similarity, and **PK-iL** with a Gaussian Kernel.

**Suicidality Context Capture** It is very interesting to note the word2vec, trained using the CBOW method, is the best performing model in the Baseline, Cosine Similarity, and the Gaussian Kernel case. We hypothesize upon inspection of the embeddings that word2vec, since trained from scratch on the suicide related post corpus captures contextual dependencies between suicidality tokens and phrases much better than LMs. LMs need to be fine-tuned on very large amounts of data to adapt against non suicidality term related contexts that they have trained on using massive corpora.

From our analysis we note that for domain specific tasks such as mental health related prediction, it is perhaps better to train contextual dependencies between words and phrases from scratch as pretrained models are already heavily biased towards the contextual dependencies on the corpora that they are trained on.

**Comparing Baselines** Across all the models we see, **PK-iL** improves upon the accuracy of the baseline models by upto almost 15% points (for Longformer). Although to confirm our statement we have to rule out effects of collecting more data, adding/deleting features etc, using neural representations and limited data alone, explicitly controlling the learned model with process knowledge shows significant performance gains.

**High Dimensional Data** Our experiments indeed show that even for domain specific corpora such as posts related to suicidality, the latent dimension of the text required to learn metric spaces is indeed very high. Improving the dimensionality shows little gain in this setting. But we hypothesize that for text from broader domains (e.g. text related to mental health in general), the dimensionality expansion will show more significant improvements.

Model	Baseline	Cosine Similarity	Gaussian Kernel
Word2Vec	75% / 69%	83% / 78%	84% / 72%
XLNET	70% / 65%	80% / 69%	84% / 71%
RoBERTa	67% / 62%	70% / 64%	71% / 62%
T5	67% / 54%	72% / 52%	75% / 64%
ERNIE	68% / 62%	75% / 69%	80% / 71%
Longformer	50% / 38%	65% / 49%	67% / 48%

Table 2: The mean accuracy/AUC-ROC, rounded up, of all the models - Column one shows the baseline where the model is directly used to predict the label, Column two shows **PK-iL** with Cosine Similarity for Kernel choice for each embedding model choice of representation, Column two shows **PK-iL** with a Gaussian Kernel for Kernel choice for each embedding model choice of representation

## 5.2 Qualitative Evaluation

As mentioned earlier the qualitative evaluations among three expert annotators received a score of 0.7 agreement. Now, We will look at some of the explanations generated for interesting examples that show cases where **PK-iL** performed well and cases where it did not. We will also compare with explanations of the word2vec model which is easy to visualize using the weights of the word2vec vectors from the logistic regression model. We highlight the phrases whose individual word sums are greater than a threshold.

**Post Example 1.** We will compare Word2vec baseline and **PK-iL** with the Gaussian Kernel.

From this example, we can clearly see word2vec associating phrases and words that characterize a low mood with suicidal ideation. In real life such words may raise triggers in the minds of a clinician and may benefit their analysis. However, the human annotator seems to have labeled this as indication based on the “there can be humor in everything” part of the post. Recall that **PK-iL** deals with whole fragments of text and can therefore never highlight phrases as we experimented with fragment lengths of 1-3 sentences. The highest threshold among the similarity functions in Equation 3 corresponded to the fragment highlighted and the path 1. Wish to be dead *no* and hence the model picks indication with probability equal to inter-annotator agreement of  $y = indication$  at that leaf. Such an explanation although subject to annotator agreements

<b>Prediction:</b> Ideation <b>Ground Truth:</b> Indication <b>Model:</b> Word2Vec Baseline
'A book is usually what I do when Im getting down, but it doesnt work when I start getting panicky. Ill try the carbs, the caffeine doesnt work because Ive gotten it in a movie theater and had a soda with me...'. 'A few reasons. I feel backed into a corner mostly. And Im Tired of being Tired of everything. If that makes sense.'. 'Thank you! I understand its a sad thing. But I also want people to realize that there can be humor in anything and its the best way to deal with this. Its how I would do it. ', 'I really dont want to ask for help. Id rather not let anyone know Im having these kind of issues.'

Table 3: Example of attention visualization based explanations

<b>Prediction:</b> Indication <b>Ground Truth:</b> Indication <b>Model:</b> PK-iL with Gaussian Kernel
'A book is usually what I do when Im getting down, but it doesnt work when I start getting panicky. Ill try the carbs, the caffeine doesnt work because Ive gotten it in a movie theater and had a soda with me...'. 'A few reasons. I feel backed into a corner mostly. And Im Tired of being Tired of everything. If that makes sense.'. 'Thank you! I understand its a sad thing. But I also want people to realize that there can be humor in anything and its the best way to deal with this. Its how I would do it. ', 'I really dont want to ask for help. Id rather not let anyone know Im having these kind of issues.'
<b>Explanation:</b> 1. Wish to be dead ( <i>no</i> ) → indication

Table 4: Example of explanation based on PK-iL

is more informative to the clinician about the models prediction. Although, the highlights from the Word2vec model provide important cues as to the user’s suicidal thought patterns it is unclear to the clinician why certain words were highlighted and certain others ignored. For example, why just “panicky” and not the whole phrase “getting panicky”?

**Embeddings vs PK-iL explanations - Developer vs End-user Perspective:** Computer scientists with deep understanding of logistic regression weights and biases may find the embedding model based visualization easier to understand and replicate. They would clear understand the contextual dependencies between tokens and phrases learned by the inner mechanism of Word2vec and could therefore reasonably

expect if the weights of a token or phrase are high in logistic regression, contextually related words will also be high. Also, that roughly statistically frequent tokens and the most frequent co-occurring words, per class label, are most likely to be highlighted in the model explanation. This makes perfect sense to the developer. However, the domain expert will struggle to palate the idea of statistically likely words and frequently co-occurring words as a valid explanation for the prediction.

**Post Example 2.** We will see PK-iL with Gaussian Kernel outputs for a slightly more interesting example.

<b>Prediction:</b> Behavior or Attempt <b>Ground Truth:</b> Behavior or Attempt <b>Model:</b> PK-iL with Gaussian Kernel
'I wish I could give a shit about what would make it to the front page. I have been there and got nothing. Same as my life. I do have a gun.'. 'I thought I was talking about it. I am not on a ledge or something, but I do have my gun in my lap.'. 'No. I made sure she got an education and she knows how to get a job. I also have recently bought her clothes to make her more attractive. She has told me she only loves me because I buy her things. '
<b>Explanation:</b> 1. Wish to be dead ( <i>yes</i> ) → 2. Non-Specific Active Suicidal Thoughts ( <i>yes</i> ) → Active Suicidal Ideation with Some Intent to Act, without Specific Plan ( <i>yes</i> ) → Behavior or Attempt

Table 5: Example of explanation based on PK-iL for example 2

Here we can see how PK-iL highlights multiple sentences that satisfy its explanation generated. Note that the Word2Vec models prediction was also correct in this instance highlighting phrases such as “On a ledge”, “have a gun” and “gun in my lap”.

**Correctness of Prediction** For the example post considered, the correctness of the prediction is subject to interpretation by human experts. This is why there is interannotator disagreement. PK-iL is however theoretically capable of performing as well as the annotators as per Assertion 3 in the best case. The intuitions behind PK-iL focus more on understanding the experts thought process and providing explanations that they can understand rather than on the correctness of prediction.

Thus we believe fundamental algorithmic and data annotation changes like the PK-iL paradigm will result in faster integration of assistive machine learning technology in real-world applications.

## 6 Conclusion and Future Work

In this study we develop a novel paradigm PK-iL that introduces the need for richer annotation and high performance explicit process guided explanation models that the end-user can readily understand. The dataset contains a lot of noisy and

long posts. In such settings both **PK-iL** and embedding models performed poorly. These inherent challenges of social media data will need to be addressed in future work. Additionally, **PK-iL** also has the potential to identify regions of the example space that the **PK** applies to with high inter-annotator agreement. This can assist in soliciting more refined guidelines on those cases where scales such as the CSSRS clearly do not work. While these scales have been developed over decades of research, machine learning techniques such as **PK-iL** have the potential to provide assistive refinement of existing and established guidelines.

## 7 Acknowledgements

This research is supported by National Science Foundation (NSF) Award # 2133842 “EAGER: Advancing Neuro-symbolic AI with Deep Knowledge-infused Learning.” Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

## References

- [Adadi and Berrada, 2018] Amina Adadi and Mohammed Berrada. Peeking inside the black-box: a survey on explainable artificial intelligence (xai). *IEEE access*, 6:52138–52160, 2018.
- [Alambo *et al.*, 2019] Amanuel Alambo, Manas Gaur, Usha Lokala, Ugur Kursuncu, Krishnaprasad Thirunarayan, Amelie Gyrard, Amit Sheth, Randon S Welton, and Jyotishman Pathak. Question answering for suicide risk assessment using reddit. In *2019 IEEE 13th International Conference on Semantic Computing (ICSC)*, pages 468–473. IEEE, 2019.
- [Beltagy *et al.*, 2020] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- [Bjureberg *et al.*, 2021] Johan Bjureberg, Marie Dahlin, Andreas Carlborg, Hanna Edberg, Axel Haglund, and Bo Runeson. Columbia-suicide severity rating scale screen version: initial screening for suicide risk in a psychiatric emergency department. *Psychological medicine*, pages 1–9, 2021.
- [Bubeck and Sellke, 2021] Sébastien Bubeck and Mark Sellke. A universal law of robustness via isoperimetry. *Advances in Neural Information Processing Systems*, 34, 2021.
- [Gaur *et al.*, 2021] Manas Gaur, Vamsi Aribandi, Amanuel Alambo, Ugur Kursuncu, Krishnaprasad Thirunarayan, Jonathan Beich, Jyotishman Pathak, and Amit Sheth. Characterization of time-variant and time-invariant assessment of suicidality on reddit using c-ssrs. *PLoS one*, 16(5):e0250448, 2021.
- [Liu *et al.*, 2019] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [Lundberg and Lee, 2017] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.
- [Mikolov *et al.*, 2013] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [Raffel *et al.*, 2019] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*, 2019.
- [Ribeiro *et al.*, 2016] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144, 2016.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [Wang *et al.*, 2018] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- [Yang *et al.*, 2019] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019.
- [Zhang *et al.*, 2019] Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. Ernie: Enhanced language representation with informative entities. *arXiv preprint arXiv:1905.07129*, 2019.