CAN NATURAL LANGUAGE PROCESSING DETECT THE POLITICS OF PAIN?

Reading Between the Lines: Using NLP to Understand Online Conversations

Samantha Butcher, Beatriz De La Iglesia

US Capitol riot, January 6: grievance reframed as resistance

Brexit campaigners claim marginalisation by political elites

INTRODUCTION

How does divisive discourse gain traction, and can machines learn to recognise its structure?

Using NLP techniques, this project investigates the language of **weaponised victimhood (WV)** in political speech.

> **Weaponised victimhood** *refers to the strategic use of pain and grievance to gain moral authority, deflect criticism, or justify action[1]*

This project uses **annotated political speeches** to train **a BERT-based NLP pipeline** combining **NER** (Named Entity Recognition) with **a positioning layer**: entities are labelled by **rhetorical role** and **assigned a value,** e.g. **victims** as **positive**, **aggressors** as **negative.**

This allows the model to track patterns of harm and blame, ***exposing who gets cast as the hero, the villain, or the victim.***
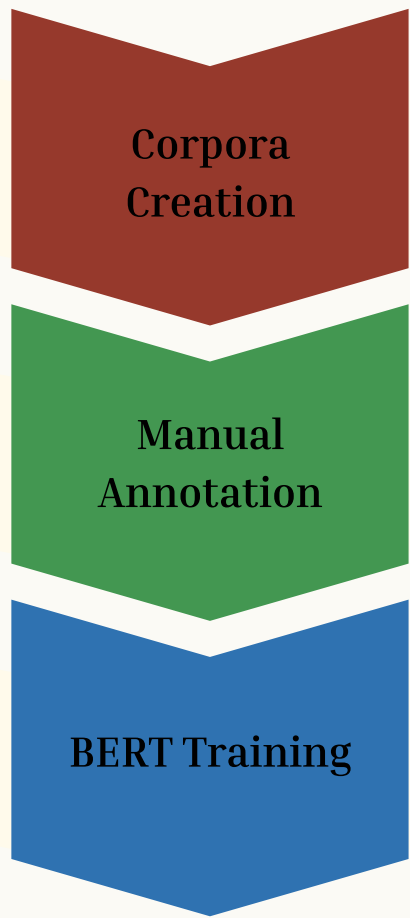

Freedom rhetoric frames public health as oppression

> "**They** want to silence **me** because **I** will never let **them** silence **you**. And in the end, **they're** not after **me**, **they're** after **you**, and **I** just happen to be standing in **their way**."
>
> Donald J. Trump, 2019

Persecutors/oppressors | Victims/innocents

METHOD

Corpora Creation
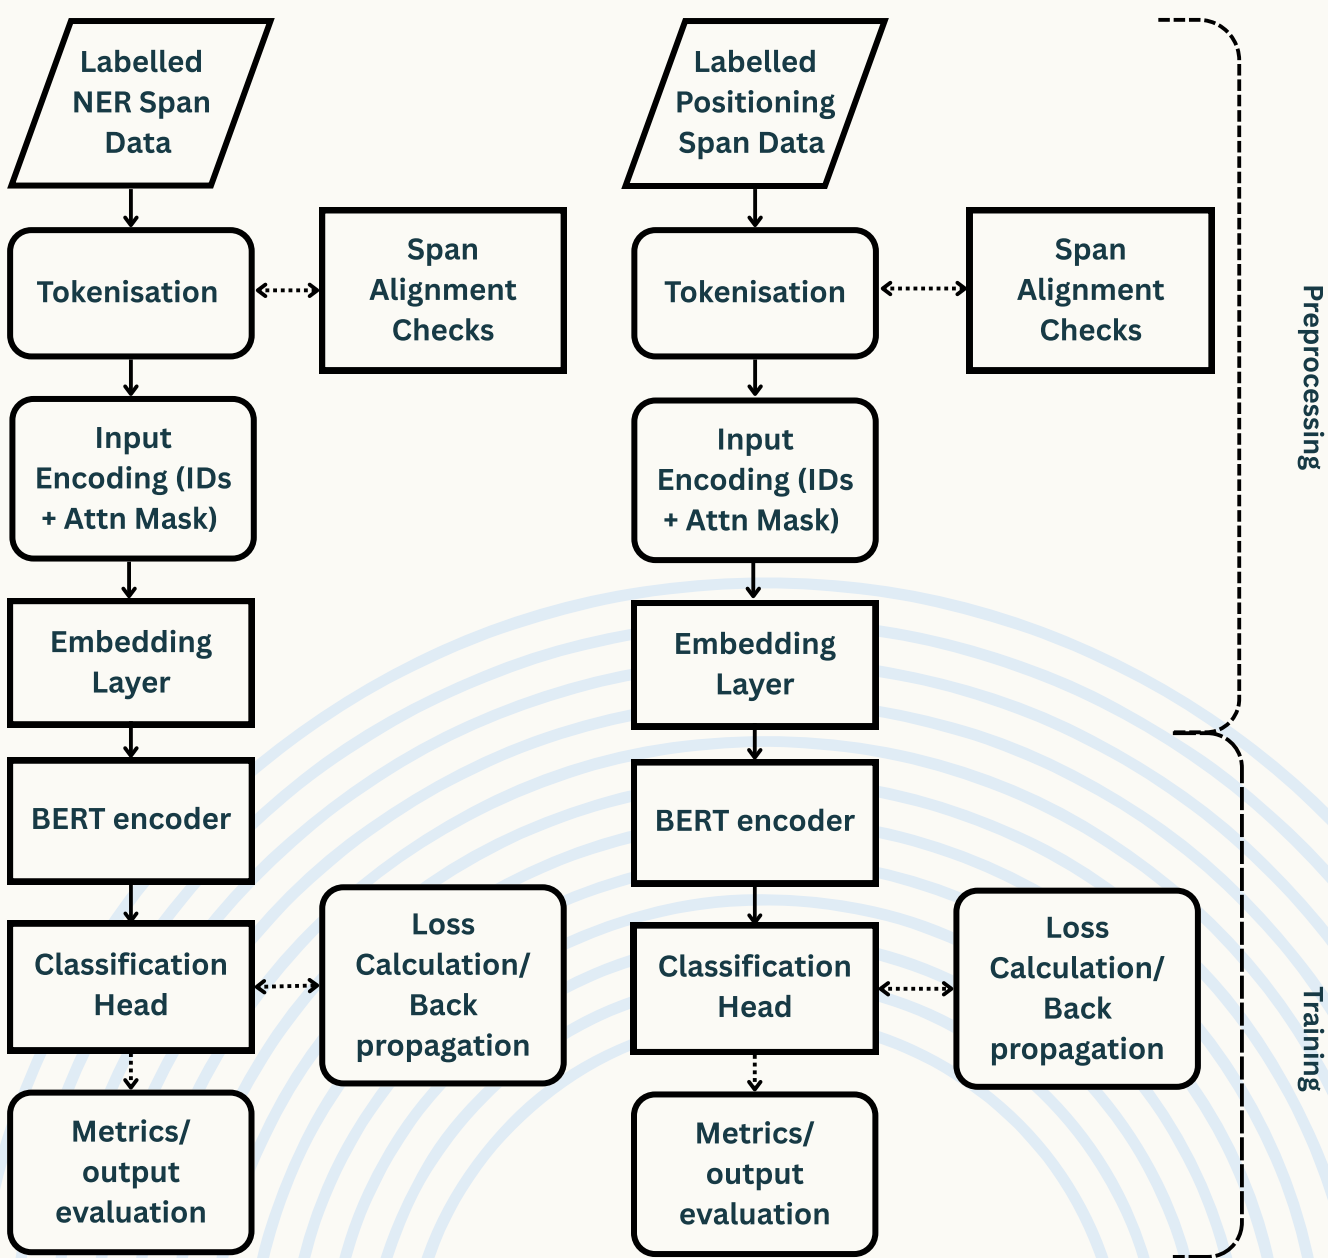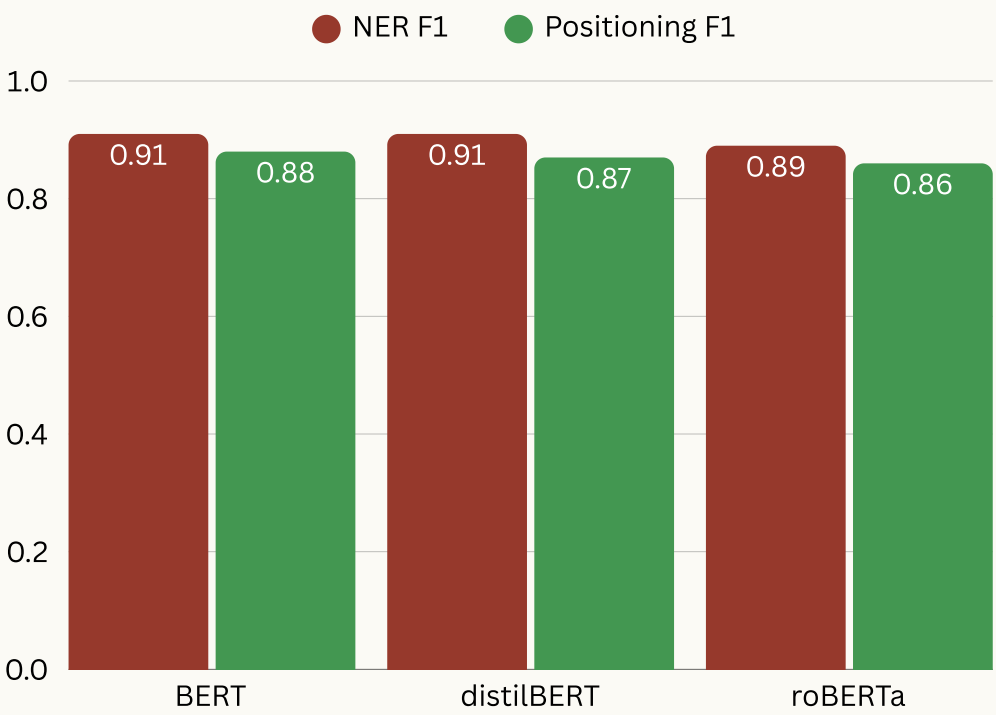
- **A curated corpus of political speeches** was compiled, focusing on emotionally charged language, group identity references, and rhetorical intensity.

Manual Annotation

- E**ach entity (you, they, us) is assigned a positioning label:** positively, neutrally or negatively positioned.

BERT Training

- **A fine-tuned NLP pipeline learns to detect these patterns** of ingroup and outgroup framing across political discourse, using the annotated labels to guide its learning.

> In the initial experimentation, three different variations of BERT were identified to test to see which performed better with these two tasks - BERT (base), distilBERT and roBERTa. distilBERT offers a lightweight version of BERT, whilst RoBERTa has optimised training and larger data exposure.

SINGLE TASK TRAINING PIPELINE

Preprocessing: Labelled NER Span Data → Tokenisation ↔ Span Alignment Checks → Input Encoding (IDs + Attn Mask) → Embedding Layer

Training: BERT encoder → Classification Head ↔ Loss Calculation/Back propagation → Metrics/output evaluation

Preprocessing: Labelled Positioning Span Data → Tokenisation ↔ Span Alignment Checks → Input Encoding (IDs + Attn Mask) → Embedding Layer

Training: BERT encoder → Classification Head ↔ Loss Calculation/Back propagation → Metrics/output evaluation

RESULTS

- NER F1
- Positioning F1

	NER F1	Positioning F1
BERT	0.91	0.88
distilBERT	0.91	0.87
roBERTa	0.89	0.86

- **NER task showed high accuracy** across all models.
- **Positioning was more challenging** due to subtle rhetorical variation and complexity.
- **BERT had best overall performance**.

FUTURE

- **Expand the corpus** with additional online sources (e.g. news articles) to **boost generalisability**.
- **Test new pipeline architectures**, including multi-task and task-conditioned learning.
- **Apply the best-performing model to Reddit data** to evaluate performance in informal, fragmented discourse, with a view to see if WV appears in other contexts.

1: Bebout, L. (2019). Weaponizing victimhood: Discourses of oppression and the maintenance of supremacy on the right. In A. Nadler & A. J. Bauer (Eds.), News on the Right (pp. 64–83). Oxford University Press.