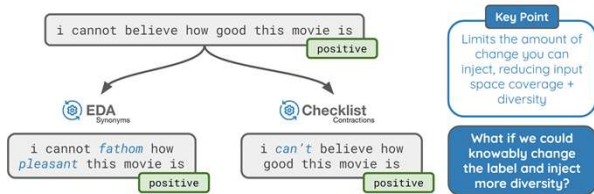


Sibylvariant Transformations for Robust Text Classification


Fabrice Harel-Canada¹, Muhammad Ali Gulzar², Nanyun Peng¹, Miryung Kim¹
UCLA¹, Virginia Tech²

Invariant Transforms


Nearly all transformations are constrained to **preserve** the source label



Unified Framework

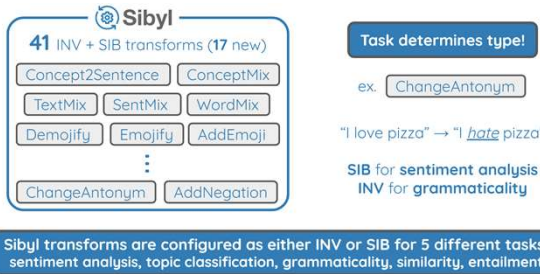



Evaluation: INV vs SIB




Sibylvariant Transforms

Jointly transform the input and label



```
> pip install sibyl-tool
```


Adaptive SIB Training



- Periodically assess model performance by class
- Generate more examples by targeting commonly confused classes
 - ex. mix "sports" topics with "politics" more often

<https://github.com/UCLA-SEAL/Sibyl>

How does SIB help?



UMAP embeddings of inputs by class

- SIB diversifies datasets more than INV to improve input space coverage
- SIB data may support margin maximizing decision surfaces