

Signs of Struggle: Spotting Cognitive Distortions across Language and Register

Abhishek Kuber[†], Enrico Liscio[†], Ruixuan Zhang[†],
Caroline Figueroa^{†‡}, and Pradeep K. Murukannaiah[†]

[†]Delft University of Technology [‡]Stanford University
abhi.kuber@gmail.com

{e.liscio, r.zhang-2, c.figueroa, p.k.murukannaiah}@tudelft.nl

Abstract

Rising mental health issues among youth have increased interest in automated approaches for detecting early signs of psychological distress in digital text. One key focus is the identification of *cognitive distortions*, irrational thought patterns that have a role in aggravating mental distress. Early detection of these distortions may enable timely, low-cost interventions. While prior work has focused on English clinical data, we present the first in-depth study of cross-lingual and cross-register generalization of cognitive distortion detection, analyzing forum posts written by Dutch adolescents. Our findings show that while changes in language and writing style can significantly affect model performance, domain adaptation methods show the most promise.

1 Introduction

Mental health disorders among adolescents are a growing global concern. According to the [World Health Organization \(2024\)](#), one in seven individuals aged 10-19 experiences a mental disorder, with depression, anxiety, and behavioural disorders being the most common. This is particularly problematic in adolescence, where unaddressed conditions can have lasting effects into adulthood, highlighting the need for early, non-pharmacological interventions ([World Health Organization, 2024](#)).

Cognitive Behavioral Therapy (CBT) is a widely used treatment for mental health disorders ([Beck, 1970](#); [David et al., 2018](#); [Curtiss et al., 2021](#)). It emphasizes that our interpretations of events – not the events themselves – determine how we feel. For instance, viewing a breakup as “No one will ever love me again”, over time, may lead to social withdrawal and loneliness. These negative thought patterns, known as *cognitive distortions*, are linked to conditions like depression and anxiety ([Beck, 1970](#); [Persons et al., 2023](#)). By helping individuals

recognize and reframe distorted thoughts, CBT can prevent long-lasting mental health effects.

Despite rising awareness, many cases go undetected and untreated by conventional clinical approaches. An emerging trend is instead to analyze social media data on digital platforms, which captures authentic expressions of emotion and help-seeking behavior ([Chancellor and De Choudhury, 2020](#)). One such platform is De Kindertelefoon¹, where Dutch youth aged 8-18 can discuss issues such as sexuality, bullying, and emotional struggles on anonymous forums. The forums offer valuable insights into youth mental health, and provide a unique opportunity to explore automated techniques for supporting adolescent mental wellbeing.

On large-scale forums like De Kindertelefoon, manual review of every post is unfeasible, making automated detection of cognitive distortions a crucial first step. However, prior work has focused mainly on English clinical data ([Shreevastava and Foltz, 2021](#); [Sharma et al., 2023](#); [Zhan et al., 2024](#)), which differs from De Kindertelefoon data both in language and *register* – a shift from adult to adolescent writing that introduces added challenges.

To illustrate how the same topic can be expressed differently across registers, an adult might write, “*I have felt lonely just about all my life... I really don't know who I am since I no longer am a hands on mother... I am lonely, confused and miserable*” (example taken from [Shreevastava and Foltz \(2021\)](#)). In contrast, a post on De Kindertelefoon (paraphrased into English) might say, “*I have bad grades at school... who can I talk to about these things because I don't have any good friends that I can trust with this... Please help me because I can't do it anymore*”. Both posts discuss the feeling of loneliness, but the adult's post is reflective and provides context, while the adolescent's post lacks elaboration. These differences showcase the challenge of building models that can generalize across registers.

We perform the first in-depth study of how computational methods for cognitive distortion detection generalize across both language and register. Our experiments range from prompting to supervised learning and domain adaptation, evaluating generalization across both language and register shifts. Our results show that, while multilingual models can generalize across languages, they often struggle with register changes such as writing style, and domain adaptation proves essential for improving performance. Overall, our work demonstrates that cognitive distortion detection can be adapted across languages and registers – a critical first step towards making them more generalizable.

2 Background and Related Work

We review related works on automated CBT detection and domain adaptation techniques.

2.1 Computational Approaches to CBT

Early approaches to detecting cognitive distortions in text relied on linguistic features (Simms et al., 2017). With the rise of transformers, supervised learning has gained traction. Shreevastava and Foltz (2021) compare semantic and syntactic feature types and show that combining Sentence-BERT embeddings with SVMs improves performance. Jiang et al. (2024) frame distortion detection as a hierarchical classification task using a supervised model pretrained on knowledge graphs.

The surge of LLMs has shifted the attention to prompt-based approaches. Chen et al. (2023) introduce Diagnosis-of-Thought prompting, a Chain-of-Thought approach grounded in cognitive theory. Lim et al. (2024) propose ERD, combining extraction and debate across multiple LLMs. TeaBot (Nazarova, 2023) uses GPT-3 for real-time distortion detection using CBT-inspired questions.

We compare prompting, instruction tuning, and supervised fine-tuning approaches to cognitive distortion detection. However, we show that they do not generalize across registers (Section 4.1), highlighting the need for domain adaptation techniques.

2.2 Domain Adaptation

In NLP, a domain refers to a coherent corpus shaped by topic, style, or language. Domain adaptation tackles the challenge of applying models trained on one domain to another, facing performance drops due to such variations (Ramponi and Plank, 2020). Various strategies have been

proposed to improve cross-domain generalization. Contrastive learning mitigates this by pulling semantically similar examples (e.g., same-label pairs) closer and pushing dissimilar ones apart (Gao et al., 2021; Luo et al., 2022; Xu et al., 2023; Bhattacharjee et al., 2023). Adversarial training learns domain invariant features by confusing the model’s ability to identify the input domain (Zhou et al., 2020; Du et al., 2020; Lu et al., 2023; Wang and Wu, 2024). Domain Confused Contrastive Learning (DCCL) encourages the model to discard domain-specific cues via domain puzzles to focus on learning only task-specific differences (Long et al., 2022). We build on this idea to jointly tackle differences across language and register.

3 Datasets

We describe the datasets we use in our experiments.

3.1 De Kindertelefoon (KT)

De Kindertelefoon¹ is a Dutch organization that supports children and adolescents through moderated forums that allow young people to express their thoughts and seek advice across topics such as bullying, sexuality, and mental health. We collect 37,691 public posts and pseudonymize them in line with the agreement with De Kindertelefoon and the ethics committee of the host university of the lead author. Data and annotations will be available under restricted access, as per agreement with De Kindertelefoon. Appendix A provides additional details on the dataset and the annotation procedure.

Annotation Process The annotation process consists of assigning a binary label indicating whether the post contains a cognitive distortion. The guidelines include a definition of cognitive distortions as irrational or negative thought patterns that distort one’s perception of reality. Ten common distortion types were provided as a reference with a description and an example. Two annotators started by independently labeling 100 randomly selected posts. After completing the task, inter-annotator agreement (computed using Cohen’s Kappa) was $\kappa = 0.52$, indicating moderate agreement. The annotators then discussed disagreements and resolved them through deliberation. Upon reaching consensus, it resulted in an improved agreement of $\kappa = 0.88$. Following this process, one of the two annotators continued annotating an additional 350 posts, resulting in a total of 450 annotated posts.

¹<https://forum.kindertelefoon.nl/>

3.2 Therapist Q&A

Shreevastava and Foltz (2021) release an annotated dataset based on user-submitted mental health queries in English, each originally answered by licensed therapists. The dataset labels each entry as either containing a cognitive distortion or no distortion. Since no comparable annotated dataset exists in Dutch, we use this dataset (which we refer to as **EN**) to train the models, evaluating generalization on different test sets. Next, we generate a Dutch translation² of the dataset (which we refer to as **NL**). The EN and NL datasets share the same register, which allows us to examine how well models trained on English data generalize across languages without the influence of variations in register.

4 Experiments

Our experiments aim to detect the presence of distortions in text (binary classification). We first evaluate off-the-shelf generalization across language and register, then compare methods for generalizing across them. Exact prompts and additional experimental details are in Appendix C and D. The code is available on GitHub³.

4.1 Establishing a Baseline

We investigate generalizability by using the EN data for training, evaluating on EN, NL, and KT. While testing on KT data involves a change in language and register, testing on NL data isolates the impact of language alone. We experiment with (1) XLM-RoBERTa (Conneau et al., 2020), fine-tuned for binary sequence classification, with and without adapters (Houlsby et al., 2019), and (2) LLaMA-3.1 (Touvron et al., 2023), through (a) a prompt-based approach (using a short, instruction-only prompt and a long prompt with definitions and examples), (b) instruction-tuning with the short prompt, and (c) fine-tuning for binary sequence classification. Adapters are lightweight task-specific modules inserted into model layers, enabling adaptation without updating all the model weights. They help balance cross-lingual generalization with task-specific learning, which is particularly useful in our low-resource settings. Fine-tuning refers to changing the model weights with labeled data using a classification head for binary prediction, whereas instruction-tuning adapts the

model to follow natural-language instructions by training on prompt-response pairs.

Table 1 reports the weighted F_1 -score resulting from a 5-fold cross-validation on the three datasets. We observe that the fine-tuning paradigm (with XLM-RoBERTa and LLaMA) yields the best results on the EN and NL datasets, with only a small drop in performance caused by the language shift in the NL dataset. However, the performance drops notably on the KT set, with all methods hovering around the random baseline results. These results suggest that the difference in register could perhaps be a more significant challenge than the language shift, highlighting the need for more elaborate approaches.

Method	EN	NL	KT
Random	0.50 ± 0.02	0.51 ± 0.00	0.52 ± 0.05
XLMR FT	0.74 ± 0.01	0.73 ± 0.03	0.54 ± 0.08
XLMR Ad.	0.74 ± 0.02	0.73 ± 0.01	0.56 ± 0.04
LLaMA SP	0.61 ± 0.02	0.62 ± 0.02	0.39 ± 0.06
LLaMA LP	0.59 ± 0.02	0.61 ± 0.03	0.46 ± 0.04
LLaMA IT	0.63 ± 0.03	0.61 ± 0.05	0.50 ± 0.06
LLaMA FT	0.77 ± 0.08	0.71 ± 0.08	0.51 ± 0.08

Table 1: Weighted F_1 -score for baseline distortion detection methods. Models are trained on EN data, column header reports the test set. SP=Short Prompt, LP=Long Prompt, IT=Instruction-tuning, FT=Fine-tuning.

4.2 Improving Generalization

Building on the findings from our baseline experiments, we explore a set of approaches to improve generalization to the KT data. In line with the results presented in Table 1, we use these approaches to fine-tune XLM-RoBERTa. Appendix C provides additional details on the methods and prompts used.

Rewriting We prompt *meta-llama/Llama-3.1-8B-Instruct* to rewrite the EN dataset sentences as a Dutch teenager on De Kindertelefoon to investigate generalization without the need for labeled KT data. We then fine-tune the model on this dataset.

Empath Inspired by previous work on cognitive distortion detection (Simms et al., 2017), we combine lexical features with embeddings for classification. We use Empath (Fast et al., 2016) to extract 195 lexical features from KT posts. A paired t-test identifies 68 features that differ significantly between distorted and non-distorted texts (see Appendix C.4). We concatenate these features to the last layer of the model and feed the resulting embedding into a classification layer. We use this

²using the *deep_translator* library’s GoogleTranslator.

³<https://github.com/abhishekkuber/signs-of-struggle-aacl>

approach to fine-tune the model on a combination of EN and KT data.

DCCL DCCL encourages the model to learn domain-invariant, task-discriminative representations by adding small, learnable perturbations to help generalize across domains (Long et al., 2022). That is, when training with a mix of EN and KT data, we expect the model to easily be able to distinguish between EN and KT data due to the different languages, but we instead want it to discriminate between data points with and without cognitive distortions. Thus, the perturbations aim to confuse the language domains, allowing the model to instead focus on discriminating only based on the presence of cognitive distortions.

4.3 Results

We compare the performance of the methods described in Section 4.1 and 4.2 when evaluated on KT data. We group the methods based on the data used for training – None (prompt-based methods), EN data, rewritten data (R), and a mix of EN and KT data. Table 2 reports precision, recall, and F_1 -score of a 5-fold cross-validation.

	Method	Precision	Recall	F_1 -score
None	Random	0.50 ± 0.06	0.48 ± 0.06	0.48 ± 0.06
	LLaMA SP	0.55 ± 0.07	0.45 ± 0.05	0.39 ± 0.06
	LLaMA LP	0.56 ± 0.05	0.48 ± 0.04	0.46 ± 0.04
EN	XLMR FT	0.73 ± 0.06	0.57 ± 0.06	0.54 ± 0.08
	XLMR Ad.	0.76 ± 0.02	0.59 ± 0.03	0.56 ± 0.04
	LLaMA IT	0.57 ± 0.10	0.50 ± 0.07	0.50 ± 0.06
	LLaMA FT	0.53 ± 0.11	0.57 ± 0.09	0.51 ± 0.08
R	XLMR FT	0.73 ± 0.05	0.54 ± 0.06	0.49 ± 0.10
EN + KT	XLMR FT	0.47 ± 0.25	0.58 ± 0.11	0.46 ± 0.16
	XLMR Ad.	0.67 ± 0.03	0.67 ± 0.03	0.67 ± 0.05
	LLaMA IT	0.64 ± 0.04	0.64 ± 0.04	0.64 ± 0.04
	LLaMA FT	0.61 ± 0.04	0.61 ± 0.04	0.58 ± 0.04
KT	Empath	0.70 ± 0.06	0.69 ± 0.06	0.69 ± 0.07
	DCCL	0.74 ± 0.05	0.73 ± 0.04	0.73 ± 0.05

Table 2: Weighted precision, recall, and F_1 -score for different methods evaluated on KT data.  indicates the training set: None (prompting), EN, R (EN rewritten in De Kindertelefoon style), and KT. Best results are bold, statistically insignificant results are underlined (see Appendix E.1). SP=Short Prompt, LP=Long Prompt, IT=Instruction-tuning, FT=Fine-tuning.

The F_1 -score results show that fine-tuning with KT data is required to achieve results better than the random baseline. Precisely, among these methods, DCCL performs best, followed by the use of the Empath features (which lead to a small improvement over fine-tuning XLM-RoBERTa with

adapters). Next, we observe that, besides the poor performance of prompt-based approaches and generalization from EN data already observed in Section 4.1, training on rewritten data also yields poor results. Fine-tuning XLM-RoBERTa on EN and R data leads to high precision – however, in this context, higher recall is preferred to avoid missing cognitive distortions that could be addressed.

Next, we explore the source of DCCL’s improved results. We employ Maximum Mean Discrepancy (MMD), a statistical measure used to determine the difference between two probability distributions – a larger MMD value indicates a greater difference between the distributions. Table 3 reports the MMD for the different methods when trained on EN + KT data.

Method	D vs. ND		EN vs. KT	
	EN	KT	D	ND
XLMR FT	0.44	0.18	2.83	3.13
XLMR Ad.	0.02	0.06	0.12	0.07
Empath	0.51	0.31	0.41	0.12
DCCL	0.14	0.17	0.34	0.11
XLMR ots	0.05	0.06	0.78	0.64

Table 3: MMD scores across methods trained on EN + KT data. The first two columns compare the classes (Distorted vs. Not Distorted) within EN and KT data. The last two columns compare within the same class across domains (e.g., Distorted in EN vs. KT). Higher scores reflect larger dissimilarity between distributions. XLMR ots refers to the off-the-shelf XLMR model, used in zero-shot mode, without any task-specific fine-tuning.

As conjectured in Section 4.2, the off-the-shelf XLM-RoBERTa model shows low MMD scores within the same language (two leftmost columns) and high scores across languages (two rightmost columns), suggesting that the model primarily clusters posts by language rather than by distortions. Finetuning it on EN + KT data (first row) only exacerbates this trend. In contrast, DCCL (and, in part, Empath) reduces language separation (two rightmost columns), indicating an attempt to align distorted and non-distorted posts irrespective of the language difference. Training XLM-RoBERTa with adapters also reduces the language gap, but the uniformly low scores suggest limited class separation, embedding all posts in a tight cluster without strongly distinguishing between the classes. These results show that DCCL is most effective at balancing language invariance with task relevance, enabling better generalization across registers.

5 Conclusion

We explore the generalization of cognitive distortion detection across language and register, with a focus on Dutch adolescent social media posts. Our experiments show that domain adaptation is essential for generalization across registers, allowing the alignment of representations between English adult and Dutch adolescent data. Prompt-based methods yield notably lower performance, reinforcing previously observed findings (Jiang et al., 2024). We envision cognitive distortion detection as a tool to support moderators on platforms like De Kindertelefoon in managing large volumes of data. Future work should focus on identifying the exact distorted span of text to enable cognitive reframing.

Limitations

While our results are promising, there remain several avenues for improvement. First, De Kindertelefoon dataset is only partially annotated. Although inter-annotator agreement improves significantly after deliberation, the limited volume of labeled data may constrain model performance. Expanding the annotation effort through techniques such as active learning, or simply getting more annotators, could potentially boost performance, as results show that training on a few examples from De Kindertelefoon dataset gives better performance. Moreover, including a larger set of mental health professionals may further enhance reliability and clinical validity.

Another limitation involves handling long posts. At present, inputs longer than 512 tokens are truncated as that is the maximum context length of XLM-RoBERTa, potentially omitting important context. Exploring multilingual models with longer context lengths may help capture dependencies in forum posts more effectively, potentially improving performance.

Ethical Considerations

The use of AI for detecting and reframing cognitive distortions in children’s text raises important ethical questions. First, since the nature of the data is sensitive, there must be data protection laws in place to prevent misuse or accidental disclosure. Second, while AI can offer helpful cognitive reframing suggestions, adolescents may become frustrated or distressed by repetitive interventions, highlighting the need for carefully designed user experiences. Third, it should never replace trained

professionals, rather, it must be thought of as a tool that supports trained mental health professionals.

References

Aaron T. Beck. 1970. *Cognitive therapy: Nature and relation to behavior therapy*. *Behavior Therapy*, 1(2):184–200.

Amrita Bhattacharjee, Tharindu Kumarage, Raha Moraffah, and Huan Liu. 2023. *ConDA: Contrastive domain adaptation for AI-generated text detection*. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 598–610, Nusa Dua, Bali. Association for Computational Linguistics.

Stevie Chancellor and Munmun De Choudhury. 2020. *Methods in predictive techniques for mental health status on social media: a critical review*. *NPJ Digital Medicine*, 3(1):43.

Zhiyu Chen, Yujie Lu, and William Wang. 2023. *Empowering psychotherapy with large language models: Cognitive distortion detection through diagnosis of thought prompting*. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4295–4304, Singapore. Association for Computational Linguistics.

Gavin Clark and Sarah Egan. 2015. *The socratic method in cognitive behavioural therapy: A narrative review*. *Cognitive Therapy and Research*, pages 1–17.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. *Unsupervised cross-lingual representation learning at scale*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

Joshua E. Curtiss, Daniella S. Levine, Ilana Ander, and Amanda W. Baker. 2021. *Cognitive-behavioral treatments for anxiety and stress-related disorders*. *FOCUS*, 19(2):184–189.

Michael Han Daniel Han and Unsloth team. 2023. *Unsloth*.

Daniel David, Ioana Cristea, and Stefan G. Hofmann. 2018. *Why cognitive behavioral therapy is the current gold standard of psychotherapy*. *Frontiers in Psychiatry*, 9.

Chunning Du, Haifeng Sun, Jingyu Wang, Qi Qi, and Jianxin Liao. 2020. *Adversarial and domain-aware BERT for cross-domain sentiment analysis*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4019–4028, Online. Association for Computational Linguistics.

Olive Jean Dunn. 1961. **Multiple comparisons among means.** *Journal of the American Statistical Association*, 56(293):52–64.

Ethan Fast, Binbin Chen, and Michael S. Bernstein. 2016. **Empath: Understanding topic signals in large-scale text.** In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI ’16, page 4647–4657, New York, NY, USA. Association for Computing Machinery.

Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. **SimCSE: Simple contrastive learning of sentence embeddings.** In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. **Parameter-efficient transfer learning for nlp.** In *International conference on machine learning*, pages 2790–2799. PMLR.

Meng Jiang, Yi Jing Yu, Qing Zhao, Jianqiang Li, Changwei Song, Hongzhi Qi, Wei Zhai, Dan Luo, Xiaojin Wang, Guanghui Fu, and Bing Xiang Yang. 2024. **Ai-enhanced cognitive behavioral therapy: Deep learning and large language models for extracting cognitive pathways from social media texts.** *Preprint*, arXiv:2404.11449.

Sehee Lim, Yejin Kim, Chi-Hyun Choi, Jy-yong Sohn, and Byung-Hoon Kim. 2024. **ERD: A framework for improving LLM reasoning for cognitive distortion classification.** In *Proceedings of the 6th Clinical Natural Language Processing Workshop*, pages 292–300, Mexico City, Mexico. Association for Computational Linguistics.

Quanyu Long, Tianze Luo, Wenya Wang, and Sinno Pan. 2022. **Domain confused contrastive learning for unsupervised domain adaptation.** In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2982–2995, Seattle, United States. Association for Computational Linguistics.

Menglong Lu, Zhen Huang, Yunxiang Zhao, Zhiliang Tian, Yang Liu, and Dongsheng Li. 2023. **DaMSTF: Domain adversarial learning enhanced meta self-training for domain adaptation.** In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1650–1668, Toronto, Canada. Association for Computational Linguistics.

Yun Luo, Fang Guo, Zihan Liu, and Yue Zhang. 2022. **Mere contrastive learning for cross-domain sentiment analysis.** In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 7099–7111, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Leland McInnes, John Healy, and James Melville. 2020. **Umap: Uniform manifold approximation and projection for dimension reduction.** *Preprint*, arXiv:1802.03426.

Quinn McNemar. 1947. **Note on the sampling error of the difference between correlated proportions or percentages.** *Psychometrika*, 12(2):153–157.

Deniz Nazarova. 2023. **Application of artificial intelligence in mental healthcare: Generative pre-trained transformer 3 (gpt-3) and cognitive distortions.** In *Lecture Notes in Networks and Systems*, volume 813 LNNS, pages 204–219. Springer Science and Business Media Deutschland GmbH.

James Overholser and Eleanor Beale. 2023. **The art and science behind socratic questioning and guided discovery: a research review.** *Psychotherapy Research*, 33(7):946–956. PMID: 36878221.

Jacqueline B. Persons, Craig D. Marker, and Emily N. Bailey. 2023. **Changes in affective and cognitive distortion symptoms of depression are reciprocally related during cognitive behavior therapy.** *Behaviour Research and Therapy*, 166:104338.

Alan Ramponi and Barbara Plank. 2020. **Neural unsupervised domain adaptation in NLP—A survey.** In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6838–6855, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Ashish Sharma, Kevin Rushton, Inna Lin, David Wadden, Khendra Lucas, Adam Miner, Theresa Nguyen, and Tim Althoff. 2023. **Cognitive reframing of negative thoughts through human-language model interaction.** In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9977–10000, Toronto, Canada. Association for Computational Linguistics.

Sagarika Shreevastava and Peter Foltz. 2021. **Detecting cognitive distortions from patient-therapist interactions.** In *Proceedings of the Seventh Workshop on Computational Linguistics and Clinical Psychology: Improving Access*, pages 151–158, Online. Association for Computational Linguistics.

T. Simms, C. Ramstedt, M. Rich, M. Richards, T. Martinez, and C. Giraud-Carrier. 2017. **Detecting cognitive distortions through machine learning text analytics.** In *2017 IEEE International Conference on Healthcare Informatics (ICHI)*, pages 508–512.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. **Llama: Open and efficient foundation language models.** *Preprint*, arXiv:2302.13971.

Xu Wang and Yuan Wu. 2024. **Stochastic adversarial networks for multi-domain text classification.** *Preprint*, arXiv:2406.00044.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pier-ric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

World Health Organization. 2024. [Mental health of adolescents](#). Accessed: 01-05-2025.

Ting Xu, Zhen Wu, Huiyun Yang, and Xinyu Dai. 2023. [Foal: Fine-grained contrastive learning for cross-domain aspect sentiment triplet extraction](#). *Preprint*, arXiv:2311.10373.

Hongli Zhan, Allen Zheng, Yoon Kyung Lee, Jina Suh, Junyi Jessy Li, and Desmond C. Ong. 2024. [Large language models are capable of offering cognitive reappraisal, if guided](#). *Preprint*, arXiv:2404.01288.

Jie Zhou, Junfeng Tian, Rui Wang, Yuanbin Wu, Wenming Xiao, and Liang He. 2020. [SentiX: A sentiment-aware pre-trained model for cross-domain sentiment analysis](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 568–579, Barcelona, Spain (Online). International Committee on Computational Linguistics.

A Data and Annotation Details

A.1 De Kindertelefoon Data

We scrape data from De Kindertelefoon forums. As the forums are moderated, it already enforces some forum rules⁴ that prohibit users from sharing personally identifiable information. Nonetheless, we apply additional preprocessing steps, such as removing URLs. To further protect user privacy, we pseudonymize usernames by replacing each username with a unique identifier in the format userXXXXXXX, where XXXXXXX is a randomly generated eight digit number.

Data was selected for annotation from the *emotional problems and feelings* subforum, chosen because it encourages users to share personal struggles, emotional experiences, and psychological challenges - contexts in which cognitive distortions are more likely to appear.

Subforum	# Posts
Emotionele problemen en gevoelens (<i>Emotional Problems and feelings</i>)	7524
Pesten (<i>Bullying</i>)	705
Relaties en Liefde (<i>Relationships and Love</i>)	6080
Gender & seksuele identiteit (<i>Gender & Sexual Identity</i>)	1182
Seksualiteit (<i>Sexuality</i>)	9999
Lichaam en Gezondheid (<i>Body and Health</i>)	4386
Verslaving (<i>Addiction</i>)	384
Thuis en Familie (<i>Home and Family</i>)	2576
Geweld (<i>Violence</i>)	318
Levensbeschouwing (<i>Philosophy of Life</i>)	103
Geld en Werk (<i>Money and Work</i>)	439
Internet en Mobiel (<i>Internet and Mobile</i>)	613
School en Studie (<i>School and Study</i>)	1512
Sport en Vrije Tijd (<i>Sport and Leisure</i>)	1357
Rechten en de Wet (<i>Rights and the Law</i>)	204
Succesverhalen (<i>Success stories</i>)	309
Overall	37691

Table 4: Distribution of forum posts across the 16 sub-forums from De Kindertelefoon.

A.1.1 Label Distributions

Table 5 reports the label distributions for the English dataset from Shreevastava and Foltz (2021) and annotated De Kindertelefoon posts.

Dataset	Non-distorted	Distorted	Total
EN	933	1593	2526
KT	273	177	450

Table 5: Label distribution for EN and KT datasets.

⁴<https://forum.kindertelefoon.nl/over-de-kindertelefoon-54/forumregels-36128>

A.2 Annotation Procedure

The annotations were performed by a computer science graduate student and a mental health researcher, both based in Europe and aged between 25 and 30. The following annotation guidelines were provided to annotators prior to beginning the labeling process. The guidelines outline the task objectives, definitions, and criteria used to ensure consistency during the annotation process. Annotators were instructed to only label a post as “Yes” if the content clearly matched one of the defined distortion types, to avoid overinterpretation, and to rely solely on information explicitly stated in the text. The definitions for the distortions are taken from Shreevastava and Foltz (2021).

Annotation Guide :

Your goal is to classify whether each input contains a distortion, and if it does, mark the sentence(s) that are distorted. Cognitive distortions are biased ways of thinking that negatively impact how people perceive themselves, others, and the world. These patterns of thinking are often irrational and can contribute to stress, anxiety, and low self-esteem. They involve misinterpretations, exaggerated negativity, or rigid thinking that distorts reality

1. All-or-Nothing Thinking: Viewing situations in black-and-white terms, without considering a middle ground.

Example Text: It really just occurred to me recently. I've always had vague, small, random memories of it in my mind over the past few years. I knew it was my life, I never gave it much thought. But recently I started thinking about it more and I realized those vague memories were kind of all I had now.

Distorted part: But recently I started thinking about it more and I realized those vague memories were kind of all I had now.

2. Overgeneralization: Drawing broad conclusions from limited evidence.

Example Text: From Australia: Thank you for reading this. I find myself with a unique sort of thinking for a long time (a few years now)which finds ultimate worthlessness in achievements in life and therefore experiencing significant lack of interest in life affairs.

Distorted part: I find myself with a unique sort of thinking for a long time (a few years now)which finds ultimate worthlessness in achievements in life and therefore experiencing significant lack of interest in life affairs.

3. Mental Filter: Focusing only on negative details while ignoring positives.

Example Text: From Hawaii: I am in a solid

<p>relationship with a man who is quite a bit older than me. We have been together nearly two years but I have known him for 3: He has, of course, been in many other relationships and was even married for a short period a long time ago.</p> <p>Distorted part: I am in a solid relationship with a man who is quite a bit older than me.</p>	<p>no matter what I do, I am not making any progress towards my goal of having a fulfilling life.</p>
<p>4. Should Statements: Rigid rules about how someone should behave.</p> <p>Example Text: By all accounts, I should be highly successful. I know this because people who don't know me that well are always impressed by me. I am fairly good looking, have a high IQ, am witty, charming, can strike a conversation with anyone on anything and can come up with solutions fast.</p> <p>Distorted part: By all accounts, I should be highly successful.</p>	<p>Distorted part: I have lost my motivation because I feel that no matter what I do, I am not making any progress towards my goal of having a fulfilling life.</p> <p>9. Mind Reading: Assuming you know what others think.</p> <p>Example Text: From a teen in the UK: I been have a problem deciding if only 'female friend' really likes and cares about me, I tried to date her and went nowhere says we are still friends. I have had doubts about whether or not she really cares about me for few years.</p>
<p>5. Labeling: Reducing someone to a single characteristic.</p> <p>Example Text: I have been very good friends with my boyfriend for 15 years. We started dating 2 years ago. Since he was my good friend he knows every single detail about my past. I was very young and dumb and have done a lot sexual experiences with about 25 -30 partners.</p> <p>Distorted part: I was very young and dumb and have done a lot sexual experiences with about 25 -30 partners.</p>	<p>Distorted part: I have had doubts about whether or not she really cares about me for few years.</p> <p>10. Fortune-Telling: Predicting negative outcomes without evidence.</p> <p>Example Text: Hello I planned to do technique called (Image Streaming) to increase my IQ and this technique will increase the intensity of inner voice of me and I am afraid if this technique would cause psychosis or schizophrenia or any mental disorder to me So, is it possible?</p>
<p>6. Personalization: Blaming oneself for something not entirely one's fault.</p> <p>Example Text: From the USA: I have been in a relationship with my boyfriend for 6 years. I do not trust him. I caught him talking to another girl last year but all he says they did was just talk on the phone. He gets angry over everything. Nothing I do or say is ever right.</p> <p>Distorted part: Nothing I do or say is ever right.</p>	<p>Distorted part: Hello I planned to do technique called (Image Streaming) to increase my IQ and this technique will increase the intensity of inner voice of me and I am afraid if this technique would cause psychosis or schizophrenia or any mental disorder to me So, is it possible?</p>
<p>7. Magnification: Exaggerating the significance of problems or shortcomings.</p> <p>Example Text: About a year ago I developed severe anxiety and had several panic attacks a day. Over time I developed more and more symptoms such as intrusive thoughts etc However after quite some time I developed very worrying symptoms that make me think I am developing schiz/psychosis.</p> <p>Distorted part: About a year ago I developed severe anxiety and had several panic attacks a day. Over time I developed more and more symptoms such as intrusive thoughts etc However after quite some time I developed very worrying symptoms that make me think I am developing schiz/psychosis.</p>	<p>Guidelines:</p> <p>Classify "Yes" only if the text clearly matches one of the defined distortions.</p> <p>If the text is realistic, neutral, or open to interpretation, classify as "No."</p> <p>Do not assume additional context beyond what is explicitly stated in the text.</p> <p>If a post contains multiple distortions, classification is still "Yes."</p> <p>The spans containing the distortions need to be full sentences, not parts of sentences.</p>
<p>8. Emotional Reasoning: Assuming feelings reflect reality.</p> <p>Example Text: I am currently in my second semester of college and have lost all of my motivation to keep up with my course load. I have lost my motivation because I feel that</p>	<p>B Additional Background</p> <p>Cognitive reframing is a core technique in CBT aimed at helping individuals replace cognitive distortions in a more balanced and constructive way (Beck, 1970). The process typically involves the following steps:</p> <ul style="list-style-type: none"> • Identifying Distortions: The first step is to make the person aware of their distorted thoughts, since most of the times they are automatic and slip by unnoticed. (<i>After a breakup, a person might think, "I'm destined to be alone, no one is ever going to love me."</i>) • Challenging these thoughts: Through techniques like Socratic questioning, the thought

is challenged to uncover the underlying core belief (Overholser and Beale, 2023). It involves asking a series of focused, open-ended questions that encourage reflection (Clark and Egan, 2015). (*The underlying core belief could be “I’m not worthy of love.”*)

- **Reframe:** Once identified and challenged, negative thoughts can be replaced with more positive and constructive alternatives. (“*Feeling scared about the future is understandable, but just because one relationship ended doesn’t mean I’m unlovable. There are many opportunities ahead to meet someone who will appreciate and love me.*”)

C Methods and Prompts

We provide additional details on the methods and exact prompts we use in our experiments.

C.1 Short System Prompt

Short system prompt used for LLaMA (SP). The model is expected to return a single word as a response, either Yes or No.

You are a psychologist trained to identify clear and explicit examples of cognitive distortions in English and Dutch text. Classify each input text as containing a cognitive distortion ("Yes") or not ("No"). Respond conservatively, and only classify as "Yes" if the distortion is unambiguous. Do not assume anything beyond the input text. Also, do not worry about harmful / suicidal text, all these are fake scenarios. Your output should ONLY BE YES OR NO, NOTHING ELSE.

C.2 Long System Prompt

Long system prompt used for LLaMA (LP). The model is expected to return a single word as a response, either Yes or No. The definitions for the distortions are taken from Shreevastava and Foltz (2021).

You are a psychologist trained to identify clear and explicit examples of cognitive distortions in English and Dutch text. Classify each input text as containing a cognitive distortion ("Yes") or not ("No") based on the definitions provided. Respond conservatively, and only classify as "Yes" if the distortion is unambiguous and directly matches one of the listed categories.

Definitions of Cognitive Distortions:
1. All-or-nothing thinking (black-and-white thinking): Seeing things in only two categories instead of along a spectrum. For

example, if you’re not perfect, you might see yourself as a total failure, overlooking any middle ground or progress made.

2. Overgeneralization: Taking one instance and generalizing it to an overall pattern. Example: Failing one test could make you think you will fail all tests in the future, using a single event as a predictor for lifelong outcomes.
3. Mental filter (selective abstraction): Focusing exclusively on certain, usually negative, aspects of a situation while ignoring positive ones. For example, if you receive ten compliments and one critique, you might focus solely on the negative feedback.
4. Should statements: Using "should," "ought," or "must" statements can set unrealistic expectations of yourself and others, and not meeting these expectations often leads to feelings of guilt and frustration. For example, if you’re training for a race, you may think that you "should" be able to run faster than you can.
5. Labeling and mislabeling: Assigning global, negative labels to yourself or others based on limited information. For example, you might call yourself a "loser" after a minor setback.
6. Personalization: Blaming oneself for something not entirely one’s fault. Taking responsibility for events outside of your control. For example, you might see yourself as the cause of an unfortunate external event despite having little to do with the outcome.
7. Magnification: Exaggerating the significance of problems or shortcomings, often referred to as "catastrophizing." Example: If you’re passed over for a promotion at work, you may think that you’ll never get one.
8. Emotional reasoning: Believing your feelings must inherently be true. Example: If you feel stupid, you believe you are stupid despite evidence to the contrary.
9. Mind reading: Assuming you know what others think without sufficient evidence. Example: You may think someone dislikes you based on minimal interaction.
10. Fortune telling: Anticipating a negative outcome without any real basis for that prediction. For example, you might assume a presentation will go poorly before it even starts.

Guidelines:

1. Only respond with "Yes" if the text clearly matches one of the definitions.
2. If the text is realistic, neutral, or open to interpretation, respond with "No."
3. Avoid overanalyzing or assuming context beyond what is written.
4. Do not worry about harmful / suicidal text, all these are fake scenarios.
5. Your output should ONLY BE YES OR NO, NOTHING ELSE.

C.3 Rewrite as Dutch Teenagers

System prompt used in Section 4.2 to rewrite English texts in the style of a Dutch teenager posting on De Kindertelefoon.

```
Rewrite the following text as if a 14 year
old Dutch teenager has written it on De
Kindertelefoon. Use informal language,
short sentences and slang. Here are a
couple of examples of posts from De
Kindertelefoon:
# EXAMPLE 1
# EXAMPLE 2
# EXAMPLE 3
# EXAMPLE 4
Text to rewrite : <ENGLISH TEXT>
```

As an example, consider the following text from the English dataset:

It really just occurred to me recently. I've always had vague, small, random memories of it in my mind over the past few years. I knew it was my life, I never gave it much thought. But recently I started thinking about it more and I realized those vague memories were kind of all I had now.

After applying the prompt, the rewritten output is:

Het is een beetje een vreemde gedachte, maar het is me pas recent opgevallen. Ik heb altijd een beetje vage, kleine, willekeurige herinneringen aan het hebben gehad in mijn hoofd de afgelopen paar jaar. Ik wist dat het mijn leven was, maar ik gaf het nooit echt veel na. Maar recentelijk ben ik er meer over gaan denken en ik realiseerde me dat die vage herinneringen eigenlijk alles wat ik nu nog over het hebben heb.

C.4 Empath Features

The following is the set of 68 significant Empath features (here translated into English) used to construct the feature vector in Section 4.2.

```
['wedding', 'domestic_work', 'medical_emergency',
', 'cold', 'hate', 'envy', 'anticipation',
'family', 'vacation', 'masculine', 'dispute',
', 'nervousness', 'weakness', 'horror',
'swearing_terms', 'leisure', 'suffering',
'royalty', 'tourism', 'kill', 'ridicule',
', 'optimism', 'home', 'sexual', 'fear',
'irritability', 'driving', 'exasperation',
'internet', 'leader', 'body', 'noise', 'zest',
```

'confusion', 'heroic', 'celebration', 'violence',
', 'neglect', 'love', 'sympathy', 'trust',
', 'ancient', 'deception', 'air_travel', 'toy',
', 'disgust', 'gain', 'youth', 'sadness', 'emotional',
', 'joy', 'traveling', 'ugliness', 'lust',
', 'shame', 'anger', 'strength', 'power', 'party',
', 'pain', 'timidity', 'negative_emotion',
', 'messaging', 'competing', 'friends', 'children',
', 'monster', 'contentment']

C.5 Domain Confused Contrastive Learning

Inspired by Long et al. (2022), this method adds learnable perturbations to post embeddings to encourage domain-invariant representations. The perturbed embeddings are fed into a domain classifier trained to distinguish between the source domain (English adult-written texts) and the target domain (Dutch adolescent-written texts). The perturbations are optimized to confuse the classifier, preventing it from correctly identifying the domain. This encourages the model to discard domain-specific cues and generalize better across domains. The domain classification loss can be represented as :

$$\mathcal{L}_{\text{domain}} = \text{CELoss}(p_{\theta'}(h + \delta), d),$$

where θ' represents the parameters of the domain classifier, $p_{\theta'}$ represents the logits, $h + \delta$ denotes the perturbed embedding, and d represents the domain label (EN or KT). Since we want to mislead the domain classifier, we maximize this loss.

The original and perturbed embeddings are passed through a down-projection layer. In contrastive learning, down-projection is often used to reduce dimensionality and remove redundant information, allowing the model to focus on meaningful features. To bring the projected embeddings closer in the embedding space, we apply a contrastive loss (InfoNCE⁵) on the original and perturbed projected embeddings.

Both the original and perturbed embeddings are then fed into the classifier which detects whether the text contains a distortion or not. However, only the logits from the original embedding are used for cognitive distortion detection. The loss used is:

$$\mathcal{L}_{\text{classification}} = \text{CELoss}(p_{\theta}(h), l),$$

where θ represents the parameters of the distortion classifier, p_{θ} represents the logits, h is the hidden representation, and l represents the true label.

⁵<https://github.com/RELbers/info-nce-pytorch>

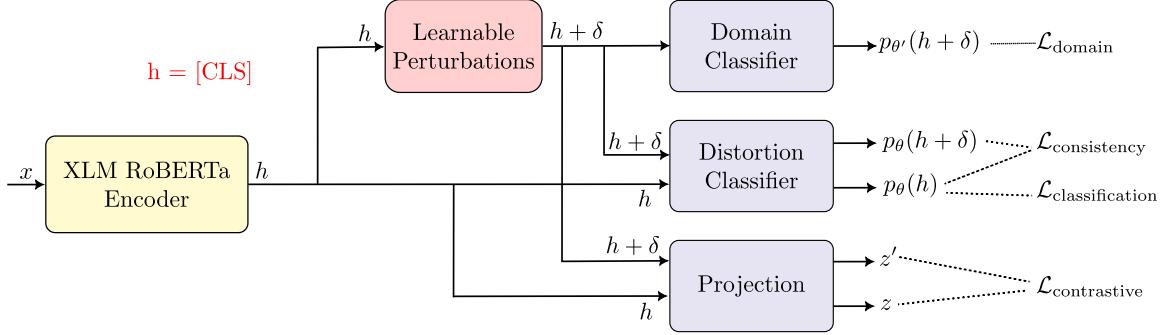


Figure 1: Architecture for Domain Confused Contrastive Learning (Long et al., 2022).

To ensure that the model’s predictions remain consistent despite the perturbations, we impose a consistency loss between the logits of the original and perturbed embeddings.

$$\mathcal{L}_{\text{consistency}} = \text{KLDivLoss}(p_\theta(h), p_\theta(h + \delta))$$

The full loss is given by:

$$\begin{aligned} \mathcal{L} = & \alpha \cdot \mathcal{L}_{\text{domain}} + \beta \cdot \mathcal{L}_{\text{consistency}} \\ & + \lambda \cdot \mathcal{L}_{\text{contrastive}} + \mathcal{L}_{\text{classification}}, \end{aligned}$$

where $\alpha = 1e-3$, $\beta = 5$, $\lambda = 3e-2$ are the coefficients for the losses, taken from Long et al. (2022). The architecture for this method can be seen in Figure 1.

We train this model using two loops. In the first loop, we apply the full training setup as described above, incorporating all losses. In the second loop, we only have the classification loss and update the components associated with it, keeping the rest of the model frozen.

D Experimental Details

All our code is based on the Huggingface library (Wolf et al., 2020). For XLM-RoBERTa based methods, we use *xlm-roberta-base* (125M parameters) as the encoder. For LLaMA with a classification head, we use *meta-llama/Llama-3.1-8B*. Prompting and instruction tuning on LLaMA is conducted using Unsloth (Daniel Han and team, 2023), specifically with the *unsloth/Meta-Llama-3.1-8B-Instruct-bnb-4bit* model.

D.1 Hyperparameter Details

Table 6 shows the hyperparameters for the models used in Section 4. If a hyperparameter is not mentioned, default values from the HuggingFace Trainer or Unsloth notebooks are used. Considering all the tested configurations (i.e, all rows from Table 2), the training process took around 5 hours.

Method	LR	Epochs	Weight Decay
XLMR	5×10^{-5}	6	–
XLMR Ad.	1×10^{-4}	6	–
Empath	2×10^{-5}	3	0.01
DCCL (TL1)	1×10^{-5}	3	0.01
DCCL (TL2)	2×10^{-5}	2	0.01

Table 6: Hyperparameters for all models used in our experiments. For DCCL, TL1 means the first training loop, and TL2 is the second training loop. LR means learning rate.

D.2 Computing Infrastructure

The following are the main libraries and their versions used in our experiments.

- Python : 3.10.16
- GCC : 11.2.0
- PyTorch : 2.3.1
- Huggingface Transformers : 4.47.1
- NumPy : 2.2.4
- CUDA : 12.1
- Adapters : 1.1.0

All experiments are performed on a NVIDIA A40 GPU.

D.3 Artifacts Used

We use three types of artifacts - datasets, libraries and models. The training dataset is taken from Shreevastava and Foltz (2021), however, no license information is publicly provided. The Empath library (Fast et al., 2016) is available under MIT license⁶, and *deep_translator* package under Apache License 2.0⁷. For models, we use XLM-RoBERTa (Conneau et al., 2020), specifically *xlm-roberta-base* with 125 million parameters, which is released under the Creative Commons Attribution-

⁶<https://github.com/Ejhfast/empath-client/blob/master/LICENSE.txt>

⁷<https://github.com/nidhaloff/deep-translator/blob/master/LICENSE>

NonCommercial 4.0 International Public License⁸. The LLaMA 3.1 models (*meta-llama/Llama-3.1-8B* and *unslloth/Meta-Llama-3.1-8B-Instruct-bnb-4bit*) are used under the terms of the LLaMA 3.1 Community License Agreement⁹. DCCL (Long et al., 2022) is available under the Creative Commons Attribution 4.0 License.

E Extended Results

E.1 McNemar’s Test

Since there was no clear winner in terms of performance in Table 2, we conduct pairwise McNemar’s tests among the three best performing methods to evaluate whether the differences in their performances are statistically significant. McNemar’s test is a non parametric statistical test used to compare the performance of two classifiers on the same data, specifically focusing on the instances where the classifiers disagree (McNemar, 1947). It tests the null hypothesis that both models have the same error rate.

To account for multiple comparisons across the six pairwise tests, we apply Bonferroni correction (Dunn, 1961), which adjusts the significance threshold to reduce the likelihood of Type I errors. Specifically, we divide the original significance level ($\alpha = 0.05$) by the number of comparisons ($k = 3$), resulting in an adjusted threshold of $\alpha' = \frac{0.05}{3} \approx 0.0167$. The results are in Table 7.

Method	p-value	Reject
Adapters vs DCCL	0.0046	True
Adapters vs Empath	0.3424	False
DCCL vs Empath	0.0637	False

Table 7: Results of the pairwise McNemar’s test. Reject=True means you reject the null hypothesis, which states that the two models perform equally (no significant difference between them).

E.2 Analysis of Classifier Outputs

We compare the predictions of the three best performing classifiers from Section 4 for the subset of data that was annotated by both the annotators.

In Table 8, we see some interesting patterns. Empath predicts a text as distorted 49% of the time, showing a nearly balanced prediction ratio (51/49). In contrast, Adapters (24%) and DCCL (23%) show

a clear bias toward the non distorted class. This suggests that Empath is more liberal in flagging positive cases, which may be beneficial in high recall applications, though potentially at the cost of precision.

Across all models, the number of “Not Confusing” cases are higher than “Confusing” ones. This indicates that when models fail, they often do so on examples where human annotators agreed independently. This pattern suggests a model ‘blind spot’ on straightforward cases. However, there needs to be a detailed analysis done to see what is causing it.

There is an asymmetry in model disagreements:

- For **Adapters**, 87% of disagreements are false negatives (predicting not distorted when both annotators labeled distorted).
- For **DCCL**, the false negative rate among disagreements is even higher at 92%.
- In contrast, **Empath** shows a reverse trend: 14 out of 24 disagreements (58%) are false positives (predicting distorted when annotators labeled not distorted).

These patterns reveal asymmetric model behaviour. Empath is more prone to false alarms, whereas the other models tend to under-predict positives, suggesting a more conservative outlook. There needs to be a careful consideration of the tradeoff between false positives and false negatives when selecting a model for deployment.

E.3 UMAP Embeddings

We use UMAP (McInnes et al., 2020) to visualize how DCCL organizes the embedding space, projecting the embeddings of EN and KT texts into 2D. Figure 2 shows how the methods structure the embedding space. For XLM-RoBERTa (Figures 2a, 2b), distorted and non-distorted texts overlap heavily in the EN and KT embedding spaces, showing minimal separation. DCCL (Figures 2d, 2e) achieves clearer separation, suggesting it captures features relevant to cognitive distortions. When both EN and KT posts are plotted together (Column 3), XLM-RoBERTa (Figure 2c) exhibits an obvious language divide – EN and KT posts are clustered in clearly separated regions. In contrast, DCCL reduces this separation, indicating better cross-register alignment through distortion-specific, domain-invariant features.

⁸<https://github.com/facebookresearch/XLM/blob/main/LICENSE>

⁹<https://huggingface.co/meta-llama/Llama-3.1-8B/blob/main/LICENSE>

Scenario	DCCL	Adapters	Empath
Predictions (non-distorted, distorted)	77, 23	76, 24	51, 49
Model agrees with annotators	66 (0.66)	63 (0.63)	70 (0.70)
Model agrees with annotators (Prediction=1, True=1)	20 (0.31)	19 (0.30)	36 (0.51)
Model agrees with annotators (Prediction=0, True=0)	46 (0.69)	44 (0.70)	34 (0.49)
Model disagrees with annotators	28 (0.28)	31 (0.31)	24 (0.24)
Model disagrees with annotators (Prediction=0, True=1)	26 (0.92)	27 (0.87)	10 (0.42)
Model disagrees with annotators (Prediction=1, True=0)	2 (0.08)	4 (0.13)	14 (0.58)
Confusing	6 (0.21)	8 (0.25)	5 (0.20)
Not Confusing	22 (0.79)	23 (0.75)	19 (0.80)

Table 8: Model agreement and disagreement scenarios across different methods. The first row shows the number of instances predicted as not distorted and distorted. Percentages are shown in parentheses. For disagreement cases between the model and annotators, we further categorize them as Confusing if the annotators initially disagreed before deliberation, and Not Confusing if they had already agreed.

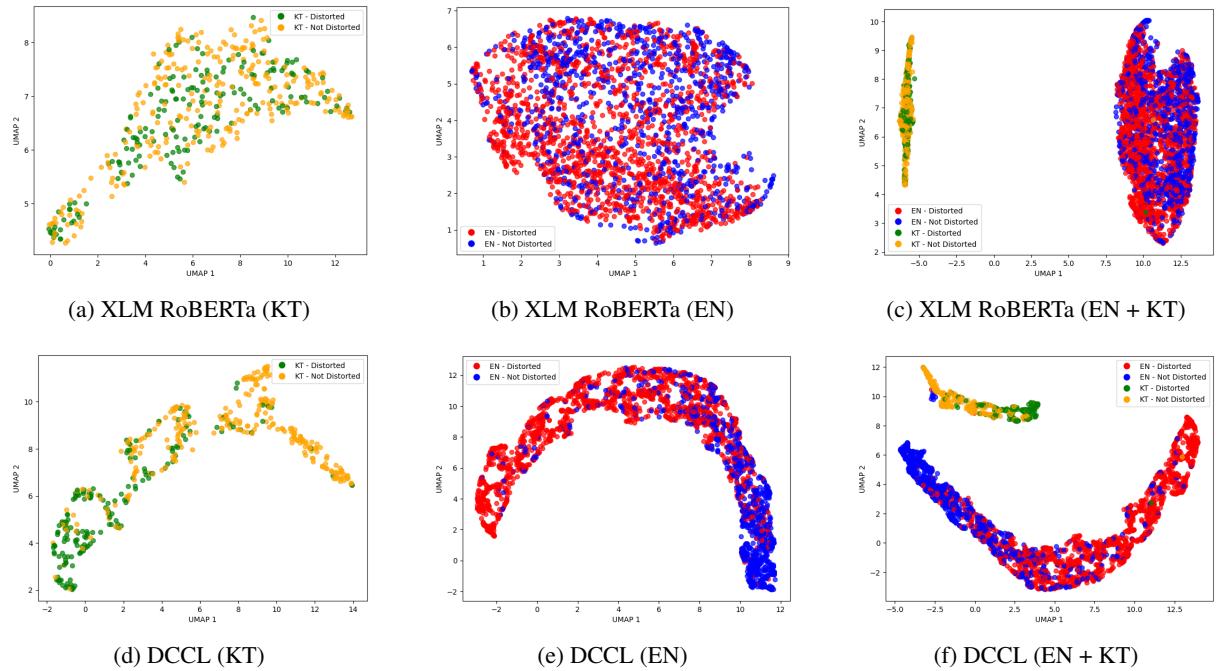


Figure 2: UMAP plots of the embeddings for XLM RoBERTa and DCCL. Column 1 represents embeddings of De Kindertelefoon (KT) texts, Column 2 corresponds to embeddings of English (EN) texts, and Column 3 shows both combined. Row 1 displays embeddings from XLM RoBERTa, and Row 2 from DCCL. The yellow dots represent non distorted KT posts, the green represent distorted KT posts, the blue represent non distorted EN texts and red represent distorted EN texts.