



Kent Academic Repository

Zhu, Zhen (2025) *Leveraging Pretrained Language Models for Maternal Health Monitoring in Online Communities*. In: Artificial Intelligence in Healthcare. Second International Conference on Artificial Intelligence in Healthcare, AliH 2025. Lecture Notes in Computer Science . pp. 75-86. Springer ISBN 978-3-032-00655-4. E-ISBN 978-3-032-00656-1.

Downloaded from

<https://kar.kent.ac.uk/111069/> The University of Kent's Academic Repository KAR

The version of record is available from

https://doi.org/doi:10.1007/978-3-032-00656-1_6

This document version

Author's Accepted Manuscript

DOI for this version

Licence for this version

UNSPECIFIED

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in **Title of Journal** , Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

Leveraging Pretrained Language Models for Maternal Health Monitoring in Online Communities

Zhen Zhu^[0000–0003–0258–1454]

Kent Business School, University of Kent, Canterbury CT2 7FS, United Kingdom
z.zhu@kent.ac.uk

Abstract. Digital maternity support communities are growing in popularity, offering valuable peer support throughout pregnancy and postpartum experiences. These platforms also generate rich textual data that can be leveraged for artificial intelligence (AI) applications. This study applies pretrained language models (PLMs) to classify and analyse 270,195 posts collected from the subreddit “BabyBumps” between 2010 and 2022. Focusing on posts that reflect personal experiences related to pregnancy, postpartum, and related events (85.9%), the analysis reveals that the majority (62.6%) centre on physical health concerns, while nearly half (48.9%) express negative sentiment. Notably, both mental health and negative sentiment-related discussions show a marked resurgence during the COVID-19 pandemic. These findings underscore the evolving emotional and informational needs of expectant and new mothers in online spaces and highlight the potential of AI-driven tools in supporting digital maternal health monitoring.

Keywords: Maternal health monitoring · Online communities · Pre-trained language models.

1 Introduction

Maternal health remains a global concern, including a range of physical and mental health challenges that affect women during pregnancy and after childbirth [30, 26, 13, 1, 18]. Complications such as gestational diabetes and hypertension often co-occur with or are compounded by mental health conditions, including prenatal and postnatal depression, anxiety disorders, and post-traumatic stress [21, 29, 31, 33]. These conditions, if untreated, can significantly impact maternal well-being, infant development, and broader family dynamics [5, 20].

Traditional support systems such as family, friends, and healthcare professionals play a crucial role in providing emotional and informational support to expectant and new mothers [28, 27, 3]. However, access to such support is not always guaranteed. Women in rural areas, those facing social isolation, or experiencing stigmatisation may lack the necessary resources or environments conducive to support-seeking [11, 4]. The need for scalable, inclusive, and responsive forms of support is therefore paramount.

In this context, digital technologies have created transformative opportunities in healthcare, particularly through the proliferation of online health communities [24]. Online maternity support communities such as “BabyBumps”, a subreddit dedicated to pregnancy, early motherhood, and related events on Reddit, have gained global traction by providing accessible spaces for peer-to-peer communication and mutual support [6, 8, 22, 23, 35]. These platforms transcend geographical, social, and temporal barriers, enabling diverse users to share experiences, exchange health information, and seek emotional reassurance in relatively anonymous and judgment-free environments [14, 15].

Alongside the rise of digital health communities, artificial intelligence (AI) has emerged as a powerful tool in modern healthcare, offering the potential for early detection, diagnosis, and personalised intervention [12, 16]. In particular, natural language processing (NLP) and machine learning have been effectively applied to analyse user-generated content across various platforms to monitor public health trends, predict mental health risks, and inform clinical decision-making [7, 19, 32, 10]. Despite growing attention to AI-based health monitoring in general online contexts, research applying such tools specifically to maternal health remains limited.

This study addresses this gap by applying zero-shot learning (ZSL) using pretrained language models (PLMs) to classify and analyse posts from “BabyBumps”. This approach enables the automated classification of posts without relying on manually annotated training data, thereby offering a scalable and efficient framework for analysing large-scale textual data. Posts are classified along multiple dimensions, including whether they are related to personal experiences, the type of health concerns (physical, mental, or none), and the sentiment expressed (positive, neutral, or negative).

The findings highlight that a considerable share of user posts express concerns related to physical health (62.6%) or negative emotions (48.9%). Notably, mental health-related posts and those with negative sentiment obtain higher community engagement, as measured by the number comments received, suggesting that users are particularly responsive to distress-related content. Temporal analysis reveals a rising trend in the volume of all types of posts, with a marked resurgence in the proportion of mental health or negative sentiment-related posts during the COVID-19 pandemic.

This study contributes to the growing body of research at the intersection of AI and maternal healthcare in at least three key ways. First, it demonstrates the feasibility and effectiveness of leveraging PLMs for automated, large-scale classification of maternal health discourse. Second, it underscores the prevalence of emotional and mental health challenges discussed within digital maternal communities especially in times of crisis, drawing attention to the importance of monitoring such issues in real time. Third, it reveals patterns of peer support and engagement that highlight the critical role these platforms play in providing social and emotional resources to expectant and new mothers.

2 Data and Methods

Data from the subreddit “BabyBumps” were collected via the Pushshift platform [2], comprising 270,195 posts submitted by 67,350 users between 2010 and 2022. For each post, the title and body text were concatenated prior to classification by a PLM.

The classification tasks performed by PLMs in this study adopt a ZSL approach, which is a learning paradigm developed to classify instances from unseen classes based on labelled examples from seen classes [34]. This is achieved by leveraging auxiliary information that relates seen and unseen classes. In text classification, ZSL allows models to categorise texts into previously unseen categories without requiring explicit training on those categories. The approach relies on embedding spaces where both text and class labels are projected. Typically, these embeddings are derived from PLMs such as Bidirectional Encoder Representations from Transformers (BERT) [9]. Let Φ represent a PLM (e.g., BERT), x be an input text, and its embedding obtained from the PLM be $\Phi(x)$.

Each class c , whether seen or unseen, is also represented in the same embedding space and is denoted as $\Phi(c)$. Then the task of text classification can be formalised as:

$$\hat{c} = \arg \max_c \text{sim}(\Phi(x), \Phi(c)) \quad (1)$$

The model predicts the unseen class with the highest similarity score, where *sim* represents a similarity measure, for instance, the cosine similarity:

$$\text{sim}(\Phi(x), \Phi(c)) = \text{cosine}(\Phi(x), \Phi(c)) = \frac{\Phi(x) \cdot \Phi(c)}{|\Phi(x)| |\Phi(c)|} \quad (2)$$

This paper employs a PLM called DeBERTa (Decoding-enhanced BERT with Disentangled Attention), a variant of BERT that introduces two advancements over BERT: a disentangled attention mechanism and an enhanced mask decoder. These innovations significantly improve the efficiency of model pretraining and the performance of downstream tasks and the model has been fine-tuned on multiple natural language inference datasets [17, 25].

The data analysed in this paper are publicly accessible and entirely observational in nature. The project received ethical approval from the Research Ethics Advisory Group at Kent Business School, University of Kent (KBSE No: 3413).

3 Results

3.1 Post Classification

Fig. 1 presents the workflow for post classification. A total of 270,195 posts submitted by 67,350 users between 2010 and 2022 were collected from the subreddit “BabyBumps”. Posts were classified along three dimensions. First, a binary classification was performed to determine whether the user was personally experiencing pregnancy, postpartum, or related events. In total, 232,202 posts (85.9%)

reflected such personal experiences and were retained for further classification, allowing the exclusion of irrelevant content such as advertisements. These personal experience posts were then further classified by the type of health concerns (physical, mental, or none) and by sentiment (positive, neutral, or negative). The results show that the majority of posts relate to physical health concerns (62.6%), and nearly half express negative sentiment (48.9%).




Fig. 1: Post classification workflow.

Fig. 2 presents the detailed distribution of sentiment (positive, neutral, and negative) across posts with different types of health concerns, i.e., no concerns, mental health concerns, and physical health concerns, in the “BabyBumps” subreddit dataset. It clearly shows that posts related to physical health concerns dominate in volume, with 145,448 posts (62.6%), while posts concerning mental health are fewer, and posts with no specific concerns are the least frequent. Among the physical health concern posts, a large proportion are marked with negative sentiment, although there is also a significant number of neutral and positive posts. This indicates that while physical issues are widely discussed, not all evoke negative emotional responses.

In contrast, posts related to mental health concerns show a strikingly high proportion of negative sentiment, making up the vast majority within that category. This is expected, as mental health-related posts are more likely to express distress, anxiety, or other negative emotions. Meanwhile, posts with no concerns show a distribution dominated by positive sentiment. This reflects the support-

ive and celebratory nature of some content shared in the community, such as announcements or milestones.




Fig. 2: Post frequency in different concerns and sentiments.

3.2 Post Content

The two bigram networks in Fig. 3 visualise the top 100 most commonly co-occurring word pairs in posts discussing physical health concerns (Fig. 3a) and mental health concerns (Fig. 3b) in the “BabyBumps” subreddit. The direction of the links in these networks represents the sequential order of the bigrams. For instance, an arrow pointing from “baby” to “shower” indicates that the bigram phrase should be read as “baby shower” rather than the reverse. In the physical health bigram network, clinical and event-based language is more commonly used. Words such as “contractions”, “cervical check”, “water broke”, “gestational diabetes”, “membrane sweep”, “early labor”, and “heart rate” indicate detailed discussions around the physiological processes of pregnancy and childbirth. Common bigrams also include references to time and progress, such as “past weeks”, “found pregnant”, and “half hour”, as well as healthcare encounters such as “appointment”, “doctor”, “went hospital”, and “sent home”. This network suggests a more informational and medically-oriented discourse, where users document symptoms, medical procedures, and timelines related to their pregnancy journeys.

On the other hand, the mental health bigram network focuses more on emotional expressions and interpersonal dynamics. The cluster centred around “feel” and words such as “need”, “vent”, “advice”, and “help” reflect personal feelings as well as support-seeking behaviour. There is also a focus on psychological struggles such as “panic attack”, “stop crying”, and “feel guilty”, highlighting the emotional vulnerability shared in these posts. The prevalence of words such as

“know people”, “family members”, and “close friends” suggests that relationships play a significant role in users’ mental wellbeing narratives. Together, these networks illustrate distinct patterns in how users discuss physical versus mental health. Similarly, a bigram network generated from posts with no concerns is available upon request.

3.3 Post Engagement

Following the classification of post types, the analysis next explores how these posts are received by the community. The number of comments serves as a useful indicator of engagement, often reflecting attention, empathy, or perceived need. Comparing the number of comments across different types of posts helps highlight where support is most actively provided and whether certain kinds of content consistently elicit more interaction. These insights can inform efforts to strengthen peer support in online communities.

The two empirical cumulative distribution function (ECDF) plots in Fig. 4 illustrate how sentiment and concern types are associated with the number of comments. In Fig. 4a, posts expressing mental health concerns receive more comments than those with no concerns or physical health concerns, particularly in the mid-range of the distribution. This indicates that users are particularly responsive to mental health-related disclosures, perhaps reflecting empathy, concern, or a recognition of shared experiences. In Fig. 4b, posts with negative sentiment consistently receive more comments than neutral or positive ones. This pattern suggests that emotionally negative content may elicit stronger responses, potentially due to the community’s motivation to offer support or advice during difficult times. All pairwise differences are statistically significant with $p < 0.01$ according to the Mann-Whitney U test.

3.4 Post Dynamics

Fig. 5 concludes the analysis by presenting the temporal dynamics of health concerns and sentiments among one-post users, as the majority of users (51.1%) posted only once in this subreddit. Both subplots use dual y-axes, with the left y-axis showing the frequency of posts in each category, and the right y-axis showing the proportion of posts specifically related to mental health concerns and negative sentiment, respectively.

Fig. 5a focuses on health concerns. Whereas all types of health concerns have increased since the subreddit’s inception in 2010, the proportion of mental health-related posts initially rises but drops around 2014, then sharply increases again from 2020 onwards. This uptick corresponds with the onset of the COVID-19 pandemic and suggests heightened psychological stress and a stronger inclination to seek support online. Fig. 5b instead focuses on sentiments. Negative sentiment consistently comprises the largest share, while all sentiment types show a clear upward trend in frequency over time. Notably, the proportion of negative sentiment remains high throughout the observed period, with a visible resurgence between 2020 and 2022, mirroring the emotional toll of the pandemic.







Fig. 3: Bigram networks.




(a) Concerns.




(b) Sentiments.

Fig. 4: Empirical cumulative distribution function.

These results demonstrate the potential of applying PLMs to large-scale online discussions. This AI-enabled approach reveals both long-term trends and crisis-induced surges in digital maternal health discourse. For example, by tracking health concerns and sentiments over time, such methods can facilitate timely responses from platform moderators and healthcare practitioners, offering a scalable solution for monitoring maternal wellbeing online.



(a) Concerns.



(b) Sentiments.

Fig. 5: Dynamics of one-post users.

4 Conclusion

This study applies zero-shot learning (ZSL) using a pretrained language model (PLM), DeBERTa, to analyse posts from “BabyBumps”, the largest online community on Reddit dedicated to pregnancy, early motherhood, and related events. By leveraging state-of-the-art artificial intelligence (AI) tools, this approach enables scalable classification of user-generated content without the need for manually labelled training data. Posts are classified along multiple dimensions, including whether the user is personally experiencing pregnancy, postpartum, or related events; the type of health concerns (physical, mental or none), and the sentiment expressed (positive, neutral, or negative).

The analysis reveals that a significant portion of posts express physical health concerns (62.6%) or negative sentiment (48.9%). Engagement analysis indicates that posts related to mental health concerns or negative sentiment receive more comments, suggesting heightened community engagement with distress-related content. Temporal analysis demonstrates an overall increase in the volume of all types of posts, with a noticeable resurgence in the proportion of posts related to mental health concerns or negative sentiment during the COVID-19 pandemic.

This study makes at least three key contributions. First, it demonstrates the effectiveness of AI tools such as PLMs for scalable textual data classification without the need for manually labelled training data, offering a method that is both efficient and adaptable to large-scale datasets. Second, it identifies key emotional and health concerns expressed in online maternal health spaces, highlighting the prevalence of mental health-related issues and their emotional impact on users, especially in times of crisis. Third, it provides insights into user engagement patterns, particularly how posts related to negative sentiment and mental health concerns consistently attract more comments, suggesting that this online community serves as a valuable space for emotional support in maternal health.

Nevertheless, this study has several limitations. First, it focuses on a single, albeit high-performing, PLM (DeBERTa) for classification tasks, which constrains the generalisability of the findings. While this choice highlights the model’s practical utility in real-world contexts, it does not offer comparative insights into other PLMs that may possess different strengths, particularly in identifying subtle or overlapping emotional and informational support signals. Future studies could systematically evaluate a wide range of PLMs under zero-shot, few-shot, and fine-tuned conditions to better understand their relative effectiveness. Moreover, integrating human validation or developing hybrid AI–human pipelines could enhance both the robustness and accountability of the classification results. These steps are critical when such results are intended to inform interventions, platform design, or clinical support systems.

Acknowledgments. This study received support from the Wellcome Trust, provided through the British Academy/Leverhulme Small Research Grant Scheme.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Alderdice, F., McNeill, J., Lynn, F.: A systematic review of systematic reviews of interventions to improve maternal mental health and well-being. *Midwifery* **29**(4), 389–399 (2013)
2. Baumgartner, J., Zannettou, S., Keegan, B., Squire, M., Blackburn, J.: The pushshift reddit dataset. In: *Proceedings of the international AAAI conference on web and social media*. vol. 14, pp. 830–839 (2020)
3. Bedaso, A., Adams, J., Peng, W., Sibbritt, D.: The relationship between social support and mental health problems during pregnancy: a systematic review and meta-analysis. *Reproductive health* **18**(1), 1–23 (2021)
4. Broadbent, R., Papadopoulos, T.: Bridging the digital divide—an australian story. *Behaviour & Information Technology* **32**(1), 4–13 (2013)
5. Burke, L.: The impact of maternal depression on familial relationships. *International review of psychiatry* **15**(3), 243–255 (2003)
6. Chivers, B.R., Garad, R.M., Boyle, J.A., Skouteris, H., Teede, H.J., Harrison, C.L.: Perinatal distress during covid-19: thematic analysis of an online parenting forum. *Journal of medical Internet research* **22**(9), e22002 (2020)
7. Coppersmith, G., Dredze, M., Harman, C.: Quantifying mental health signals in twitter. In: *Proceedings of the workshop on computational linguistics and clinical psychology: From linguistic signal to clinical reality*. pp. 51–60 (2014)
8. Denton, L.K., Creeley, C.E., Stavola, B., Hall, K., Foltz, B.D.: An analysis of online pregnancy message boards: Mother-to-mother advice on medication use. *Women and Birth* **33**(1), e48–e58 (2020)
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Burstein, J., Doran, C., Solorio, T. (eds.) *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019)
10. Di Cara, N.H., Maggio, V., Davis, O.S., Haworth, C.M.: Methodologies for monitoring mental health on twitter: systematic review. *Journal of Medical Internet Research* **25**, e42734 (2023)
11. Drentea, P., Moren-Cross, J.L.: Social capital and social support on the web: the case of an internet mother site. *Sociology of health & illness* **27**(7), 920–943 (2005)
12. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., Dean, J.: A guide to deep learning in healthcare. *Nature medicine* **25**(1), 24–29 (2019)
13. Figueiredo, B., Conde, A.: Anxiety and depression in women and men from early pregnancy to 3-months postpartum. *Archives of women's mental health* **14**, 247–255 (2011)
14. Gleeson, D.M., Craswell, A., Jones, C.M.: Women's use of social networking sites related to childbearing: An integrative review. *Women and Birth* **32**(4), 294–302 (2019)
15. Gui, X., Chen, Y., Kou, Y., Pine, K., Chen, Y.: Investigating support seeking from peers for pregnancy in online health communities. *Proceedings of the ACM on Human-Computer Interaction* **1**(CSCW), 1–19 (2017)
16. He, J., Baxter, S.L., Xu, J., Xu, J., Zhou, X., Zhang, K.: The practical implementation of artificial intelligence technologies in medicine. *Nature medicine* **25**(1), 30–36 (2019)

17. He, P., Liu, X., Gao, J., Chen, W.: DeBERTa: Decoding-enhanced BERT with disentangled attention. In: International Conference on Learning Representations (2021)
18. Hermann, A., Fitelson, E.M., Bergink, V.: Meeting maternal mental health needs during the covid-19 pandemic. *JAMA psychiatry* **78**(2), 123–124 (2021)
19. Hinduja, S., Afrin, M., Mistry, S., Krishna, A.: Machine learning-based proactive social-sensor service for mental health monitoring using twitter data. *International Journal of Information Management Data Insights* **2**(2), 100113 (2022)
20. Howell, E.A., Mora, P., Leventhal, H.: Correlates of early postpartum depressive symptoms. *Maternal and child health journal* **10**, 149–157 (2006)
21. Jiang, L., Tang, K., Magee, L.A., von Dadelszen, P., Ekeroma, A., Li, X., Zhang, E., Bhutta, Z.A.: A global view of hypertensive disorders and diabetes mellitus during pregnancy. *Nature Reviews Endocrinology* **18**(12), 760–775 (2022)
22. Jiang, L., Zhu, Z.: Information exchange and multiple peer groups: A natural experiment in an online community. *Journal of Economic Behavior & Organization* **203**, 543–562 (2022)
23. Jiang, L., Zhu, Z.: Maternal mental health and social support from online communities during pregnancy. *Health & Social Care in the Community* **30**(6), e6332–e6344 (2022)
24. Jin, C., Zhu, Z.: Multimorbidity patterns and early signals of diabetes in online communities. *JAMIA Open* **8**(3), ooaf049 (2025)
25. Laurer, M., Van Attevelde, W., Casas, A., Welbers, K.: Less annotating, more classifying: Addressing the data scarcity issue of supervised machine learning with deep transfer learning and BERT-NLI. *Political Analysis* **32**(1), 84–100 (2024)
26. Lee, A.M., Lam, S.K., Lau, S.M.S.M., Chong, C.S.Y., Chui, H.W., Fong, D.Y.T.: Prevalence, course, and risk factors for antenatal anxiety and depression. *Obstetrics & Gynecology* **110**(5), 1102–1112 (2007)
27. Orr, S.T.: Social support and pregnancy outcome: a review of the literature. *Clinical obstetrics and gynecology* **47**(4), 842–855 (2004)
28. Orr, S.T., Miller, C.A.: Unintended pregnancy and the psychosocial well-being of pregnant women. *Women's Health Issues* **7**(1), 38–46 (1997)
29. Rahman, A., Surkan, P.J., Cayetano, C.E., Rwagatare, P., Dickson, K.E.: Grand challenges: integrating maternal mental health into maternal and child health programmes. *PLoS medicine* **10**(5), e1001442 (2013)
30. Ross, L.E., McLean, L.M., Psych, C.: Anxiety disorders during pregnancy and the postpartum period: a systematic review. *Depression* **6**(9), 1–14 (2006)
31. Souza, J.P., Day, L.T., Rezende-Gomes, A.C., Zhang, J., Mori, R., Baguiya, A., Jayaratne, K., Osoti, A., Vogel, J.P., Campbell, O., et al.: A global analysis of the determinants of maternal health and transitions in maternal mortality. *The Lancet Global Health* **12**(2), e306–e316 (2024)
32. Tommasel, A., Diaz-Pace, A., Godoy, D., Rodriguez, J.M.: Tracking the evolution of crisis processes and mental health on social media during the covid-19 pandemic. *Behaviour & Information Technology* **41**(16), 3450–3469 (2022)
33. Wachs, T.D., Black, M.M., Engle, P.L.: Maternal depression: a global threat to children's health, development, and behavior and to human rights. *Child development perspectives* **3**(1), 51–59 (2009)
34. Wang, W., Zheng, V.W., Yu, H., Miao, C.: A survey of zero-shot learning: Settings, methods, and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)* **10**(2), 1–37 (2019)
35. Zhu, Z.: Maternal mental health monitoring in an online community: a natural language processing approach. *Behaviour & Information Technology* pp. 1–10 (2024)