Toward Personalized Meditation Recommendations for Digital Well-Being

Angelo Geninatti Cossatin[1,†], Noemi Mauro[1,†]

[1]University of Turin, Corso Svizzera 185, Torino, 10148, Italy

Abstract

Digital meditation platforms are widely used to support mental well-being, yet many still rely on generic recommendations that overlook users' emotional states and situational context. In this work, we explore the use of interpretable machine learning to generate personalized meditation suggestions. We frame the task as a classification problem that predicts fine-grained meditation styles from textual descriptions of user prompts and contextual information, enriched with emotion labels. Two models, Logistic Regression and Random Forest, are evaluated using embedding representations of the input text. Experimental results indicate that user prompts and contextual information provide the most informative signals for identifying suitable meditation techniques, whereas emotion labels alone contribute limited predictive value. To enhance transparency, we also perform token-level attribution analysis, highlighting how specific words and phrases influence model predictions and revealing interpretable connections between user language, emotional cues, and recommended meditation practices.

Keywords

Meditation recommendation, Emotion-aware recommender, Mental wellness

1. Introduction

Personalized and adaptive systems increasingly shape how people access health, education, media, and emotional support. Within this landscape, recommender systems have the potential to contribute not only to engagement and efficiency, but also to social good, for instance by supporting users' mental well-being. Mental health and stress management have become pressing societal concerns, with a growing body of evidence linking regular meditation practice to reductions in anxiety, improved emotional regulation, and better sleep quality [1, 2]. Digital meditation platforms (e.g. PetitBambou[1], Insight Timer[2]) have proliferated in response, yet most continue to rely on static, one-size-fits-all session recommendations that ignore the user's momentary emotional state. A user experiencing acute anger may benefit most from a different practice than one seeking to build self-confidence, yet current systems rarely distinguish between such cases in a transparent, explainable way. Some platforms do provide some form of tailored recommendation (e.g. Headspace[3], Wysa[4]) but this feature is often paywalled and opaque.

Recommender systems in adjacent domains (music, news, e-learning) have demonstrated that personalisation improves engagement and outcome. Emotion-aware personalisation adds a further dimension by conditioning recommendations on affective signals extracted from free text, physiological data, or self-reports [3]. Within the wellness domain, however, the literature on interpretable, emotion-driven meditation recommenders remains sparse. Existing commercial systems are largely opaque: users cannot verify why a particular session was suggested, which undermines trust and limits the capacity for self-directed practice.

Joint Proceedings of the ACM UMAP Workshops 2026, UMAP 2026, June 8–11, 2026, Gothenburg, Sweden

[†]These authors contributed equally.

angelo.geninatticossatin@unito.it (A. Geninatti Cossatin); noemi.mauro@unito.it (N. Mauro)
0009-0007-5378-7061 (A. Geninatti Cossatin); 0000-0001-8234-3266 (N. Mauro)

[1]https://www.petitbambou.com/
[2]https://insighttimer.com/
[3]https://www.headspace.com/
[4]https://www.wysa.com/

In line with the goals of the UMAP4Good workshop, which explores how adaptive systems can support positive societal impact, this paper investigates how simple and interpretable machine learning approaches can be used to personalize meditation recommendations based on users' emotional states and context. Our focus is on building a system that is not only effective but also transparent and understandable to users.

The paper makes these contributions:

1. We formulate meditation recommendation as a classification task based on users' emotions and context.
2. We compare two classifiers, showing that linear models offer strong performance for this task.

Together, these components form a transparent recommender system that can justify each recommendation in human-interpretable terms for trustworthy mental-wellness AI. The analysis further confirms that both classifiers achieve very good accuracy on the fine-grained meditation types. Code and data are available at https://github.com/angelogeninatti/UMAP4GoodMeditation

2. Related Work

Emotion-aware recommender systems have been widely studied as a way to improve personalization by incorporating users' affective states into recommendation models. Early work by Tkalčič et al. [3, 4] demonstrated that affective context can enhance recommendation performance beyond traditional preference-based approaches. Moreover, studies showed that modeling users' emotional states can improve recommendation quality [5], particularly in domains such as music where consumption is strongly influenced by mood [6, 7]. Subsequent research expanded these ideas by proposing architectures and algorithms that integrate emotional profiles directly into recommendation pipelines, enabling systems to adapt suggestions based on contextual affective information [8, 9].

Conversational recommender systems represent a complementary direction for delivering personalized recommendations through interactive dialogue with users [10]. Moreover, the integration of affective signals within conversational systems has been suggested as a promising approach for supporting emotionally sensitive applications such as mental health and digital well-being, where users' momentary states and preferences may change over time [11, 12]. In the context of meditation and mindfulness applications, conversational recommendation mechanisms could help bridge the gap between emotion detection and actionable guidance by interactively mapping users' reported or inferred emotional states to suitable meditation practices.

Advances in affective natural language processing have further supported this line of research by enabling more precise emotion detection from textual data. Large-scale annotated datasets such as EmoBank and GoEmotions provide rich resources for modeling emotional content in language, supporting both dimensional and categorical emotion representations [13, 14]. These datasets allow models to capture nuanced emotional signals beyond simple positive-negative sentiment and have been widely used for training emotion classification systems capable of extracting affective features from user-generated text. Such capabilities are important in recommendation scenarios where user preferences are expressed implicitly through text, reviews, or descriptions.

Despite the expanding literature on emotion-aware recommendation and explainable AI, approaches specifically targeting meditation remain relatively limited. Reviews of personalization in digital mental health applications suggest that personalization strategies can improve engagement and relevance, but the empirical evidence remains fragmented and inconsistent across studies [15]. Within the mindfulness and meditation domain, existing studies have primarily focused on evaluating the effectiveness of meditation apps or analyzing user feedback through techniques such as sentiment analysis and thematic analysis [16, 17].

Consequently, there remains limited research examining how emotional signals extracted from user-generated text can be mapped to meditation styles through transparent recommendation models. In particular, few studies have explored the combination of emotion-aware recommendation with

Table 1
Emotion distribution across the enrinched Meditation-miniset-v0.2 dataset.

Emotion	Count	Emotion	Count
curiosity	14,586	nervousness	156
caring	1,666	sadness	149
desire	1,326	fear	94
neutral	1,055	admiration	71
disappointment	309	love	39
confusion	294	optimism	31
gratitude	224		

Table 2
Sample entries from the enriched dataset.

User experience level	Context	User prompt	Suggested techniques	Meditation style	Emotion	Fine-grained meditation style
beginner	User is preparing for a creativity-enhancing session and requires detailed guidance on dealing with negative body image.	I'm struggling with negative body image. How can meditation help me accept myself?	body scan, grounding exercises, gentle breathing	mindfulness exercise	confusion	Compassionate Body Scan Meditation
advanced	User is preparing for a session focused on overcoming fear and requires guidance on maintaining emotional stability.	How can I stay centered when I feel pulled in different directions?	body scan, grounding exercises, gentle breathing	mindfulness exercise	neutral	Color Visualization Meditation
intermediate	User is preparing for a session focused on cultivating self-compassion and emotional openness.	I want to bring more playfulness into my life. Can meditation guide me?	body scan, grounding exercises, gentle breathing	mindfulness exercise	desire	Yoga Nidra Meditation

systematic explainability analysis in the context of meditation or digital well-being applications. The present study addresses this gap by proposing a recommendation framework for meditation content.

3. Dataset and Preprocessing

3.1. Source Dataset

We use the `BuildaByte/Meditation-miniset-v0.2` dataset distributed through HuggingFace [18]. The dataset contains 20,000 synthetic meditation sessions. Each record includes a user prompt, contextual information, a target meditation style, suggested techniques, intended outcome, and additional metadata describing the generated session. All meditation sessions have a duration between 15 and 20 minutes. The dataset is balanced with respect to meditation experience level: 33.5% of sessions are designed for beginners, 33.3% for intermediate practitioners, and 33.2% for advanced users.

The original dataset defines the field meditation style with only three possible values: *mindfulness exercise*, *guided meditation*, and *positive affirmation session*. To obtain a more detailed categorization, we incorporate the Kaggle Meditation Techniques Dataset[5], a curated catalog containing 68 meditation techniques. Each entry in this catalog includes a *Name*, a *Description* (a short sentence describing the purpose of the technique), and an *Instructions* field explaining how to practise it.

We use this catalog to derive a fine-grained meditation style label. Specifically, we compute TF-IDF (Term Frequency–Inverse Document Frequency) embeddings for two textual components: (i) the concatenation of the `Description` and `Instructions` fields for each technique from this catalog, and (ii) the `user prompt` from the original dataset. Cosine similarity is then computed between the prompt embedding and each technique embedding. The technique with the highest similarity score is assigned as the `fine-grained meditation style`. Using this procedure, 38 of the 68 catalog techniques receive at least one assignment.

To further enrich the dataset with affective information, we apply `SamLowe/roberta-base-go_emotions` [14], a RoBERTa-base model fine-tuned on the GoE-

[5]https://www.kaggle.com/datasets/suraj520/meditation-techniques-dataset

Table 3
Classification performance on the fine-grained meditation style. Three configurations are compared: UC, emotion only, and their concatenation (combined).

Model	Acc	F1	Prec	Rec
Logistic Regression (UC)	1.0000	1.0000	1.0000	1.0000
Random Forest (UC)	0.9263	0.9220	0.9571	0.9263
Logistic Regression (emotion only)	0.1684	0.0980	0.0948	0.1684
Random Forest (emotion only)	0.1684	0.0980	0.0948	0.1684
Logistic Regression (combined)	1.0000	1.0000	1.0000	1.0000
Random Forest (combined)	0.8895	0.8733	0.9065	0.8895

motions corpus. The model assigns one of 28 fine-grained emotion labels to each user prompt. The predicted label is stored in the new field `emotion`. This approach enables the identification of contextual emotional cues that may not be captured by simple keyword-based methods. Across the dataset, 13 of the 28 possible GoEmotions categories are observed. Table 1 reports the distribution of emotion labels in the dataset.

Table 2 shows representative examples from the enriched dataset.

4. Methodology

Given the enriched dataset, we train two classifiers to recommend the fine-grained meditation style.

1. **Logistic Regression (LR)**: an ℓ_2-regularised multinomial logistic regression model. The learned coefficients directly correspond to per-class feature weights, enabling straightforward keyword attribution.
2. **Random Forest (RF)**: an ensemble model consisting of 100 decision trees.

Both classifiers use TF-IDF features embeddings derived from the concatenation of user prompt and `context` (denoted as UC).

To evaluate the discriminative contribution of emotional information alone, we train two additional *emotion-only* variants of each classifier. In this configuration, the input consists exclusively of a 13-dimensional one-hot encoding representing the GoEmotions `emotion` label.

Finally, we train *combined* variants of LR and RF that concatenate the TF-IDF embeddings of user prompt and `context` with the one-hot encoded emotion vector.

5. Results

5.1. Model Classification Performance

Table 3 reports the classification performance for predicting the fine-grained meditation style. Three feature configurations are evaluated: UC text features, emotion-only features, and their concatenation (combined).

Logistic Regression trained on UC features achieves perfect performance across all evaluation metrics. In contrast, the Random Forest model achieves a lower performance (F1 = 0.9220), with the largest degradation observed for the following techniques: *Morning Meditation* (F1 = 0.33, recall = 0.20), *Spiritual Meditation* (F1 = 0.50, precision = 0.36), *Transcendental Meditation* (F1 = 0.57, recall = 0.40), and *Breath and Loving-Kindness in Nature Meditation* (F1 = 0.57).

The emotion-only models confirm that the detected emotional state alone provides insufficient information for recommendation (F1 = 0.0980, accuracy = 0.1684 for both LR and RF). These results indicate that textual features derived from the user prompt and context are the primary source of discriminative signal.

Table 4
Top-1 fine-grained meditation style predictions (LR) for sample queries.

Query (truncated)	Predicted fine-grained meditation style	Conf.
I'm struggling with imposter syndrome. Can meditation help me feel more confident?	Affirmation Meditation	0.994
I'm feeling emotionally numb. Can you guide me to reconnect with my emotions?	Body Awareness & Gratitude Med.	0.985
Can meditation help me develop a more positive relationship with my body?	Breath & Body Awareness w/ Affirmations	0.995
I want to open my heart to love and connection.	Breath & Loving-Kindness in Community	0.995
I want to sleep peacefully tonight. Can you help?	Evening Meditation	0.998
I feel grateful for everything in my life.	Gratitude Meditation	0.997

The combined (UC + emotion) models augment the textual features with emotion labels. Logistic Regression continues to achieve perfect performance, suggesting that the additional emotion features neither degrade nor significantly enhance the predictive capability of the linear model. In contrast, the combined Random Forest model (F1 = 0.8733) performs worse than the RF trained solely on UC features (F1 = 0.9220).

Overall, the results indicate that UC textual features remain the dominant source of predictive information. Emotion labels provide complementary affective context but do not independently drive the recommendation of fine-grained meditation techniques.

Example outputs from the best-performing model (Logistic Regression) are presented in Table 4. Confidence values are from the softmax probability of the top-ranked class.

5.2. Token Attribution Analysis

To understand the decision process of the LR model, we compute per-token attribution scores for each input.[6] For a given input text and predicted class c, each token's contribution is defined as the element-wise product of its TF-IDF weight x_i and the corresponding Logistic Regression coefficient $w_{c,i}$:

$$\text{score}(i, c) = x_i \cdot w_{c,i}$$

A positive score indicates that the token supports the prediction of class c, whereas a negative score indicates a counteracting influence.

To analyze how emotional states influence vocabulary patterns associated with specific fine-grained meditation style, we group the analysis according to the GoEmotions label of the user prompt. For each emotion category, we randomly select a test sample, apply the LR model to the combined user prompt and context input, and report the highest-scoring tokens associated with the predicted fine-grained meditation style.

Table 5 reports the three highest-scoring tokens for selected emotion–fine-grained meditation style pairs.

[6]We use scikit-learn's default English tokenizer, which removes punctuation and stopwords.

Table 5
Token attention attribution for selected GoEmotions–fine-grained meditation style pairs.

Emotion → Fine-grained meditation style	**Token**	**Score**
gratitude → Gratitude Med.	feel grateful	+1.47
	grateful life	+1.47
	grateful	+1.17
love → Loving-Kindness (Community)	heart love	+1.38
	love connection	+1.38
	open heart	+1.38
desire → Evening Med.	sleep peacefully	+1.68
	peacefully tonight	+1.68
	sleep	+1.68
confusion → Comp. Body Scan	body image	+0.87
	negative body	+0.87
	help accept	+0.87
disappointment → Loving-Kindness Med.	feelings inadequacy	+1.53
	struggling feelings	+1.53
	inadequacy	+1.53

The attribution results reveal semantically coherent relationships between user language and recommended fine-grained meditation style. Prompts labeled with *gratitude* activate tokens such as *feel grateful* and *grateful life*, which strongly support the prediction of Gratitude Meditation. Similarly, prompts associated with *love* highlight tokens such as *open heart* and *love connection*, guiding the model toward Loving-Kindness practices oriented toward interpersonal compassion.

For the emotion *desire*, sleep-related tokens (e.g., *sleep peacefully*, *peacefully tonight*) strongly activate the Evening Meditation class, indicating that temporal and physiological cues remain highly informative within TF-IDF representations. Prompts labeled as *confusion* emphasize body-related vocabulary (e.g., *negative body*, *help accept*), which aligns with the focus of Compassionate Body Scan techniques on non-judgmental somatic awareness. Finally, prompts associated with *disappointment* activate tokens related to self-evaluation and struggle, directing the classifier toward Loving-Kindness meditation practices that encourage self-compassion.

6. Conclusion and Future Work

This work explored the feasibility of building recommender system for personalized meditation guidance. We framed the task as a meditation classification problem based on user context (UC) text and emotion annotations. Experimental results show that textual UC features provide the strongest discriminative signal for identifying appropriate meditation techniques.

Overall, the findings suggest that user-described context and intent are the primary drivers of effective meditation recommendations. Emotion labels, while valuable for interpretability and contextual understanding, did not improve predictive performance in the current feature formulation.

Several directions remain for future work. First, richer affective representations could be explored, such as continuous emotion embeddings, multi-label emotional states, or transformer-based emotion detection, which may capture subtler affective nuances than the current discrete emotion features. Second, expanding the dataset with data collected from real users would improve the robustness and generalizability of the models. Third, contextual user attributes such as session goals, or temporal factors (e.g., time of day or stress level) could be incorporated to enable deeper personalization. Finally, future systems should move beyond classification toward more complex recommendation techniques. Evaluating such systems through user studies would help assess their impact on engagement, and mental well-being outcomes.

Declaration on Generative AI

During the preparation of this work, the authors used Claude[7] to: Paraphrase and reword, and Grammar and spelling check. After using these tools, the author reviewed and edited the content as needed. The author takes full responsibility for the publication's content.

References

[1] G. Chételat, A. Lutz, E. Arenaza-Urquijo, F. Collette, O. Klimecki, N. Marchant, Why could meditation practice help promote mental health and well-being in aging?, Alzheimer's Research & Therapy 10 (2018) 57. URL: https://doi.org/10.1186/s13195-018-0388-5. doi:10.1186/s13195-018-0388-5.

[2] J. Lynch, L. Prihodova, P. J. Dunne, Á. Carroll, C. Walsh, G. McMahon, B. White, Mantra meditation for mental health in the general population: A systematic review, European Journal of Integrative Medicine 23 (2018) 101–108. URL: https://www.sciencedirect.com/science/article/pii/S1876382018304591. doi:10.1016/j.eujim.2018.09.010.

[3] M. Tkalcic, U. Burnik, A. Košir, Using affective parameters in a content-based recommender system for images, User Modeling and User-Adapted Interaction (2013).

[4] M. Tkalčič, U. Burnik, A. Odić, A. Košir, J. Tasič, Emotion-aware recommender systems – a framework and a case study, in: S. Markovski, M. Gusev (Eds.), ICT Innovations 2012, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 141–150.

[5] Y. Zheng, B. Mobasher, R. Burke, Emotions in Context-Aware Recommender Systems, Springer International Publishing, Cham, 2016, pp. 311–326. URL: https://doi.org/10.1007/978-3-319-31413-6_15. doi:10.1007/978-3-319-31413-6_15.

[6] W. Assunção, L. S. G. Piccolo, L. Zaina, Considering emotions and contextual factors in music recommendation: A systematic literature review, Multimedia Tools and Applications 81 (2022) 8367–8407. doi:10.1007/s11042-022-12110-z.

[7] V. Moscato, A. Picariello, G. Sperlí, An emotional recommender system for music, IEEE Intelligent Systems 36 (2021) 57–68. doi:10.1109/mis.2020.3026000.

[8] F. Narducci, M. de Gemmis, P. Lops, A general architecture for an emotion-aware content-based recommender system, in: Proceedings of the 3rd Workshop on Emotions and Personality in Personalized Systems 2015, EMPIRE '15, Association for Computing Machinery, New York, NY, USA, 2015, p. 3–6. URL: https://doi.org/10.1145/2809643.2809648. doi:10.1145/2809643.2809648.

[9] M. Polignano, F. Narducci, M. de Gemmis, G. Semeraro, Towards emotion-aware recommender systems: an affective coherence model based on emotion-driven behaviors, Expert Systems with Applications 170 (2021) 114382. URL: https://www.sciencedirect.com/science/article/pii/S0957417420310575. doi:https://doi.org/10.1016/j.eswa.2020.114382.

[10] D. Jannach, A. Manzoor, W. Cai, L. Chen, A survey on conversational recommender systems, ACM Comput. Surv. 54 (2021). URL: https://doi.org/10.1145/3453154. doi:10.1145/3453154.

[11] H. Rashkin, E. M. Smith, M. Li, Y.-L. Boureau, Towards empathetic open-domain conversation models: A new benchmark and dataset, in: A. Korhonen, D. Traum, L. Màrquez (Eds.), Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 5370–5381. URL: https://aclanthology.org/P19-1534/. doi:10.18653/v1/P19-1534.

[12] S. M. Kambare, K. Jain, I. Kale, V. Kumbhare, S. Lohote, S. Lonare, Design and evaluation of an ai-powered conversational agent for personalized mental health support and intervention (mindbot), in: 2024 International Conference on Sustainable Communication Networks and Application (ICSCNA), 2024, pp. 1394–1402. doi:10.1109/ICSCNA63714.2024.10863855.

[13] S. Buechel, U. Hahn, EmoBank: Studying the impact of annotation perspective and representation

[7] https://claude.ai/

format on dimensional emotion analysis, in: M. Lapata, P. Blunsom, A. Koller (Eds.), Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers, Association for Computational Linguistics, Valencia, Spain, 2017, pp. 578–585. URL: https://aclanthology.org/E17-2092/.

[14] D. Demszky, D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, S. Ravi, GoEmotions: A dataset of fine-grained emotions, in: D. Jurafsky, J. Chai, N. Schluter, J. Tetreault (Eds.), Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Online, 2020, pp. 4040–4054. URL: https://aclanthology.org/2020.acl-main.372/. doi:10.18653/v1/2020.acl-main.372.

[15] P. Matthews, C. Rhodes-Maquaire, Personalisation and recommendation for mental health apps: A scoping review, Behaviour & Information Technology 44 (2025) 2389–2404. doi:10.1080/0144929X.2024.2356630.

[16] Éva Gál, S. Ștefan, I. A. Cristea, The efficacy of mindfulness meditation apps in enhancing users' well-being and mental health related outcomes: a meta-analysis of randomized controlled trials, Journal of Affective Disorders 279 (2021) 131–142. URL: https://www.sciencedirect.com/science/article/pii/S0165032720328317. doi:https://doi.org/10.1016/j.jad.2020.09.134.

[17] I. Plaza, M. M. P. Demarzo, P. Herrera-Mercadal, J. García-Campayo, Mindfulness-based mobile applications: Literature review and analysis of current features, JMIR mHealth and uHealth 1 (2013). URL: https://www.sciencedirect.com/science/article/pii/S229152221300020X. doi:https://doi.org/10.2196/mhealth.2733.

[18] BuildaByte, Meditation-miniset-v0.2 dataset, HuggingFace Datasets, 2024. URL: https://huggingface.co/datasets/BuildaByte/Meditation-miniset-v0.2.