ROMANIAN JOURNAL OF INFORMATION SCIENCE AND TECHNOLOGY
Volume 29, Number 4, 2026, 325–336
doi: 10.59277/ROMJIST.2026.4.04

Modeling Fragrance Discovery through Multi-Attribute Fusion of Olfactory Accords and Affective Profiles

Elena-Ruxandra LUŢAN[1,*] and Costin BĂDICĂ[1]

[1]Department of Computers and Information Technology, University of Craiova, Craiova, Romania
Email: elena.lutan@edu.ucv.ro*, costin.badica@edu.ucv.ro
* Corresponding author

Abstract. This paper presents a multi-dimensional model for fragrance discovery, designed to bridge the gap between crowd-sourced olfactory and affective profiles. Traditional Information Retrieval in the olfactory domain often suffers from a vocabulary mismatch, where chemical descriptors (accords) do not align with how users experience scent emotionally. The proposed model addresses this gap by integrating fragrance accords and emotional profiles extracted from user reviews into a unified multi-dimensional representation. Multi-dimensional similarity is computed using domain-appropriate measures: Fuzzy Jaccard Similarity for accord matching (accommodating the inherent vagueness of olfactory perception) and Cosine Similarity for affective alignment. Unlike deterministic systems, a probabilistic re-ranking strategy is employed to ensure discovery and diversity in search results. The system is validated using a stratified k-fold cross-validation, measuring the alignment between retrieval candidates and user preferences across scent, emotion and lifestyle characteristics. The experimental results demonstrate that integrating affective perceptions with olfactory profile significantly improves the precision of intent-based fragrance retrieval.

Key-words: Affective computing; fuzzy representation; olfactory informatics.

1. Introduction

Perfume discovery is a sensory-consumer decision problem in which choices are driven not only by product composition but also by affective and contextual meaning.

A perfume is a mixture of aromatic compounds, used to impart a pleasant smell to the human body [1] or to clothing [2], depending on where it is applied. The quality of a perfume can be assessed from at least two different and quite subjective perspectives: the view from a perfumer and the view from an end-user representing the final customer. A perfumer will be interested on the quality of the ingredients, their purity, how do they blend to produce a complex accord or how does this smell evolve over time. On the other hand, for most customers, the most significant decision factor in appreciating a perfume is simply whether they enjoy its fragrance.

The reasons for appreciating a certain scent may include a preference for specific accords in the scent composition [3] or contextual scent harmonization [4]. Beyond structural composition, scent appreciation is deeply tied to the psychological dimensions of fragrance choice and personal identity [5], which manifest distinct preferences across international mentalities [6]. At a neurological level, these preferences are reinforced by the direct link between olfactory cues and the evocation of certain emotions and memories [7]. Capturing and quantifying these latent reasons remains a significant challenge [8].

The subjective or private signature scent associated to a specific person is a scent carrying out a personal connection which brings comfort [9] and familiarity to his or her peers and it is often associated with the living presence of the individual. In this context, people tend to explore perfumes from a rather narrowed spectrum, to well maintain their sense of trust, security and belonging [10], olfactory aesthetic experiences [11], while avoiding scents that might trigger negative emotional states such as anxiety or depression [12].

Existing recommendation models for fragrance typically follow a binary approach: they either analyze the perfume as a static list of ingredients, or they treat user perception as a black-box sentiment. This leaves a significant gap in representing the scent's identity, which is neither purely chemical nor purely psychological.

This paper proposes a multi-dimensional model that fuses crowd-sourced olfactory accords with affective profiles. Unlike standard content-based systems, this approach integrates multiple attribute dimensions (olfactory accords and emotional profiles) extracted from user social data through dimension-specific similarity computation. The system utilizes natural language processing to extract emotional preferences from unstructured user reviews. By fusing olfactory profile (perfume data) with emotional associations (reviews), it identifies candidates that align with both the scent and affective profile. The retrieval quality is enhanced through a probabilistic selection model, which promotes result diversification and encourages the discovery of novel fragrances within user's preferences. The system is experimentally validated using a dataset about real perfumes collected from *Fragrantica* website – an online perfume encyclopedia. The dataset was manually collected for academic research purposes in February 2024, and consists in 1500 perfumes and associated reviews spanning October 2018 to February 2024.

The study is guided by the following research questions:

RQ1: Does joint modeling of olfactory profile and emotional profile improve fragrance retrieval compared with single-dimension retrieval?

RQ2: How sensitive is retrieval quality to the accord-emotion fusion?

RQ3: Although lifestyle characteristics are excluded from ranking, do retrieved perfumes align with contextual suitability (season, time of day, longevity, sillage, gender)?

The contributions can be summarized as follows:

1. Proposal of a multi-dimensional model for perfume representation that integrates olfactory profile (accords), affective dimensions (emotions) and contextual lifestyle. The system enables fragrance discovery through dual-matching of olfactory similarity and user emotional intent, while using contextual alignment to evaluate retrieval quality.

2. Introduction of a probabilistic re-ranking strategy that balances relevance with diversity, promoting the discovery of novel fragrances within user preference spaces, while retrieving a

single most-similar perfume per query. This single perfume approach aligns with the sequential nature of perfume discovery, where users explore one fragrance at a time, grounded in the need for focused olfactory evaluation. The probabilistic re-ranking approach provides controlled variance across repeated interactions - similar users may retrieve different perfumes from the high-similarity candidate pool, promoting exploration without overwhelming choice.

3. Experimental validation across five weight w configurations demonstrates that Precision@1 and Recall@1 remain remarkably stable across varying accord-emotion fusion ratios. This resilience reveals that olfactory profile and emotional characterization are intrinsically aligned representations of perfume identity, validating the dual-dimensional modeling approach.

4. Introduction of *Multi-Dimensional Match Rate*, a novel evaluation measure for evaluating retrieval quality across olfactory (ACCS), affective (EMS), and contextual (LCS) dimensions simultaneously, offering comprehensive assessment beyond traditional precision-based metrics. Critically, MDMR incorporates lifestyle characteristics as an external validation dimension (withheld from the retrieval algorithm) to verify whether accord-emotion-based discovery implicitly captures contextual appropriateness (seasonality, longevity etc.). The high alignment across this multi-faceted constraint confirms the system's capacity for discovery, demonstrating that it retrieves contextually suitable perfumes without explicit lifestyle-based ranking.

The paper is structured as follows. Section 2 presents related work. Section 3 describes the proposed system, which includes the gathering of experimental dataset, its preprocessing and the method used for providing recommendations. Section 4 provides an overview of the experimental dataset, introduction of the experimental setup and then discussion the experimental results. The last section presents the conclusions and points to future work.

2. Related Work

It is essential to distinguish between the various descriptors utilized in fragrance literature. *Ingredients* refer to the specific chemical compounds (natural or synthetic) that form the building blocks of a fragrance. *Notes* are the individual scents (e.g. rose, jasmine, sandalwood) perceived at different stages of evaporation (top, middle, base). *Accords* — the focus of this work — are complex, harmonious combinations of multiple notes that create a unified, distinct character. Finally, *Fragrance Families* (e.g. floral, woody, oriental) represent the high-level taxonomic categories used to classify perfumes based on their primary olfactory themes.

Research at the molecular level, such as the framework by Sharma et al. [13], uses graph-based generative models to predict odor labels from chemical structures, allowing recommendation systems to operate on a fundamental level. At a higher level of abstraction, fragrance families and notes are frequently used. Hanafizadeh et al. [14] propose an expert system utilizing 11 variables, including fragrance family and prestige, while Pimentel et al. [15] explore cross-modal correspondences between musical modes and fragrance families, demonstrating how multisensory data enhances consumer engagement. Furthermore, Tobing et al. [16] employ AHP and TOPSIS for multi-criteria recommendation (quality, price, aroma, durability), though they note that deterministic ranking often limits recommendation diversity.

Capturing the emotional impact of scents is vital for personalization. Porcherot et al. [17] developed the Geneva Emotion and Odor Scale (GEOS) to map odor-elicited emotions onto affective dimensions, a methodology refined by Jaeger et al. [18] using CATA/RATA formats to profile product-focused emotion. In recommendation systems, Toumy et al. [19] leverage sentiment analysis of user reviews to identify similar users. When reviews for new products are unavailable, they suggest calculating similarity based on notes and seasonality.

While existing methods often operate either at the molecular level (ingredients) or the taxonomic level (families or notes), the proposed model leverages accords. In this paradigm, accords are posited as the optimal bridge between raw chemical composition and subjective human perception. By integrating these accords with affective emotional profiles extracted from social data, the system moves beyond content-based systems to address the gap between structured olfactory representations and subjective affective intent.

3. System Design

The proposed fragrance discovery system integrates olfactory profile with affective signals from user reviews. The architecture is illustrated in Fig. 1 in the supplementary document of the paper [20]. Three distinct parts of the system can be identified: (1) the collection of perfumes dataset and its preprocessing, (2) the input of new user reviews and their preprocessing, and (3) the application of the proposed fragrance discovery method.

The perfumes dataset is collected from *Fragrantica* website in February 2024, for purely academic research purposes, using manual scraping of public perfume pages and associated user reviews, and *Beautiful Soup* and *Selenium* Python libraries to parse the HTML code and extract two high-dimensional entities: Perfume and Perfume Review. The attributes for each entity are presented in Tables 1 and 2 of the supplementary document of the paper [20].

In what follows the data model of the proposed system is introduced. The perfume space $\mathscr{P}$ refers to the set of perfumes. The accords space $\mathscr{F}$ refers to the *Fragrantica* set of accords present in the dataset.

A perfume $P \in \mathscr{P}$ is defined as a multi-dimensional triple

$$P = (ACC, EM, LC) \tag{1}$$

ACC refers to the crowd-sourced olfactory representation (perfume accords). Specifically, it consists of a set of *Fragrantica* established olfactory themes, where the accords prominence is determined by aggregated user-votes on the platform. Each perfume accord $A_i \in \mathscr{F}$ is present in a certain ratio a_i into the final fragrance profile. The accords are modeled as the following fuzzy set with support set $\mathscr{F}$ (Eq. 2), to account for the inherent vagueness and subjectivity of scent perception:

$$ACC = \{(A_1, a_1), (A_2, a_2), ..., (A_n, a_n)\}, A_i \in \mathscr{F}, a_i \in [0, 1] \tag{2}$$

The membership function $\mu_{ACC}(A_i) = a_i$ denotes the degree to which accord A_i is present in the perfume scent and is determined by normalizing $weight_i$ corresponding to the "Accords" field collected from *Fragrantica* website.

Emotions space E refers to the following set of W.G. Parrott primary set of emotions:

$$E = \{anger, fear, joy, love, sadness, surprise\} \tag{3}$$

EM refers to the perfume affective representation represented as the following fuzzy set with support set E:

$$\begin{aligned} EM = \{&(anger, \mu_{EM}(anger)), (fear, \mu_{EM}(fear)), \\ &(joy, \mu_{EM}(joy)), (love, \mu_{EM}(love)), \\ &(sadness, \mu_{EM}(sadness)), (surprise, \mu_{EM}(surprise))\} \end{aligned} \tag{4}$$

For each emotion $x \in E$, the membership function $\mu_{EM}(x) \in [0,1]$ is determined by aggregating the dominant emotions detected across the social review corpus for each perfume. The membership function $\mu_{EM}(x)$ is computed as the ratio of the frequency of emotion x to the total frequency of all detected emotions.

EM is build through a three-stage pipeline consisting in: *(1) Preprocessing*: Each review is normalized via lower-casing, punctuation handling, tokenization, and stop-word removal. *(2) Emotion Classification*: a Sequential Neural Network (3 dense layers of 12, 8, 6 neurons with ReLU/Softmax activations) is applied to map text to the six W.G. Parrott primary emotions. Following former approach in [21], this model is trained on the Saravia et al. [22] corpus (20k labeled phrases, 10 epochs, standard split), with labels aligned to the system's six emotional categories: anger, fear, joy, love, sadness, surprise, which are consistent with the W.G. Parrott primary emotions representation. *(3) Aggregation and Normalization*: The frequency of each emotion identified across the perfume reviews is aggregated and normalized to obtain $\mu_{EM}(x)$.

LC refers to the Lifestyle characteristics. LC is defined as a tuple of fuzzy sets:

$$LC = (SS, TD, LG, SL, GD) \tag{5}$$

where:

- SS is a fuzzy set capturing perfume's suitability for a season. It has support set $\{winter, spring, summer, fall\}$
- TD is a fuzzy set with support set $\{day, night\}$
- LG is a fuzzy set with support set $\{veryweak, weak, moderate, longlasting, eternal\}$
- SL is a fuzzy set with support set $\{intimate, moderate, strong, enormous\}$
- GD is a fuzzy set with support set $\{female, morefemale, unisex, moremale, male\}$

The membership functions denote the degree to which each lifestyle attribute belong to a specific fuzzy category: suitability for a season (SS), appropriate time of day (TD), scent longevity (SL), scent sillage (SS) and targeted gender (GD). Each membership function is calculated by normalizing the raw vote count for each support set category against the total votes received for that lifestyle attribute, mapping the distribution into the unit interval $[0, 1]$.

The first part of the proposed system involves collecting the primary data from *Fragrantica* website and mapping it onto the proposed data model.

The second part consists in receiving a user review and analyzing the perfume preferences of the user of interest on basis of the review.

The last part of the system is a multi-dimensional similarity aggregation strategy (Algorithm 1 of the supplementary material [20]). Unlike keyword-based search, this approach aligns the user's intent with both olfactory profile and emotional profiles of candidate perfumes.

Importantly, the retrieval model uses only accords (ACC) and emotions (EM) to retrieve similar perfumes. Lifestyle characteristics (LC) serve exclusively for external validation, assessing whether the retrieved perfumes also align with the held-out lifestyle attributes. The fusion of accords (ACC) and emotions (EM) is performed through the following combination of similarities:

$$\begin{aligned} \text{sim}(P_x, P_y) = w \cdot \text{sim}(ACC_x, ACC_y) + (1 - w) \cdot \text{sim}(EM_x, EM_y), \\ \text{sim}(P_x, P_y) \in [0, 1] \end{aligned} \tag{6}$$

where:

- $sim(ACC_x, ACC_y)$ is the Jaccard similarity between the accords fuzzy vectors ACC_x and ACC_y for the reviewed perfume P_x and candidate recommendation P_y. Jaccard Similarity is utilized for accords to emphasize the olfactory accords overlap.

- $sim(EM_x, EM_y)$ is the Cosine similarity between the emotional representation fuzzy vectors EM_x and EM_y are for the reviewed perfume P_x and candidate recommendation P_y. Cosine similarity is utilized to capture the directional alignment of emotional profiles.

- w is the weighting factor to control the relative importance of the accords and emotional vectors representations, $w \in [0, 1]$.

A perfume is considered a candidate, if it partially shares accords and emotions with the fragrance previously reviewed by the target user. The accords-based filtering helps maintain a sense of familiarity by preserving the comfort associated with known scent profiles, while emotions-based filtering ensures that candidate perfumes resonate with the user's affective preferences, promoting a personalized emotional connection.

A probabilistic re-ranking strategy is utilized, motivated by the necessity to balance recommendation precision with serendipitous exploration. Standard top-k ranking strategies often prioritize the most similar items, leading to a repetitive experience that lacks exploration, while the probabilistic re-ranking strategy mitigates the filter-bubble effect. Each candidate's selection probability is proportional with its recommendation score, determined by the fusion of olfactory and emotional profiles.

The system incorporates sentiment-based retrieval adaptation, i.e. dominant emotion detected in the user review guides the candidate selection. Specifically, based on two identified emotional categories (positive vs. negative), the system prioritizes either high-scoring or low-scoring items, aligning candidates with the inferred emotional context.

If the user of interest writes a positive review for a perfume P_x (i.e. the identified emotion is joy, love or surprise), then the user will be suggested a perfume which is similar to P_x. In this case, the candidates with higher score shall have a higher probability of being retrieved.

On the other hand, if the user of interest writes a negative review (i.e. the identified emotion is anger, fear or sadness), the system shall suggest a perfume which has the similar accords with P_x, but differs from P_x either by the accords' ratios or by the emotional representation EM. In this case, the candidates with lowest score shall have a higher probability of being retrieved.

Finally, the system selects a single perfume from the candidate pool using roulette wheel selection, where selection probability is proportional to similarity score.

4. Experiments and Discussion

4.1. Dataset overview

The experimental dataset consists in 1500 perfumes released across years 2019-2023 and their attributes collected from *Fragrantica*. To ensure a representative sample of the olfactory market, the top 300 perfumes per year are collected based on a composite popularity metric. This metric, as defined by domain experts in the *Fragrantica* community, reflects a multi-dimensional measure of user engagement, longevity, and historical legacy.

The perfumes dataset was collected in February 2024, for purely academic research purposes, using manual scraping of public perfume pages and associated user reviews. Because the data originates from a third-party platform, we cannot redistribute raw data in full.

Within this corpus, a total of 73 unique accords is identified. Top 40 accords are illustrated in Fig. 2 of the supplementary material [20]. This distribution confirms the dataset's coverage of the four primary Scent Families: Woody (e.g. "woody", "warm spicy", "patchouli"), Floral (e.g. "white floral", "fruity", "rose"), Amber (e.g. "vanilla", "amber"), Fresh (e.g. "citrus". "fresh", "green"), , ensuring that the retrieval framework is tested against a diverse semantic space.

For each perfume in the experimental dataset, the complete corpus of associated social reviews is retrieved as of January 2024, totaling 140677 reviews. On average, there are 93 reviews per perfume, the maximum number of reviews available for a perfume being 210. Table 3 of the supplementary material [20] shows the number of perfumes for each number interval of reviews number. The perfumes are fairly evenly distributed across the review count ranges and the significant proportion (37.07%) receiving over 100 reviews indicates a high level of user engagement. This high volume of social data is critical for constructing stable emotional profiles and overcoming the sparsity issues common in standard content-based systems.

After applying the emotion extraction method, it is observed a predominance of positive emotions (80.72% joy, love, surprise) over negative ones (19.28% anger, fear, sadness). As summarized in Table 4 of the supplementary material [20], this skew reflects a common reporting bias, where users are more inclined to document satisfying experiences. Furthermore, analysis of review length (Table 5 of the supplementary material [20]) reveals that positive reviews are generally more descriptive, with higher word counts than negative counterparts. When combined with the selection bias inherent in the popularity-based sampling, these factors confirm that the dataset represents mainstream, high-engagement fragrance discourse. While this skew limits representativeness for niche or critique-focused scent discussions, it provides a stable, high-volume foundation for modeling affective-olfactory alignments in mainstream perfume discovery.

4.2. Experimental setup

To rigorously assess the proposed system, a stratified k-fold cross-validation (k=5) is implemented. For each iteration, the review corpus was partitioned into 80% training and 20% testing. The training set is used for defining the perfume emotions for the perfumes space, while the testing set is used to simulate users writing reviews and seeking for perfume discovery.

A critical challenge in olfactory data is the uneven distribution of social engagement across perfumes. To ensure every perfume maintains a robust emotional representation while providing sufficient testing queries, the 80/20 split was applied per perfume.

The details regarding number of training and testing reviews used in each iteration, together with the corresponding emotion, are presented in Table 6 of the supplementary material [20].

In what follows, the parameters for the evaluation system are defined:

- *User space* U is the set of users.
- *Training Reviews space* TrR refers to the total number of reviews *r* which are used for defining the W.G. Parrott emotional perfume representation EM.
- *Testing Reviews space* $TstR$ refers to the total number of reviews r_x given as input and seeking for perfume discovery

$$r_x = (user_i, P_j, em_k) \tag{7}$$

where: $user_i \in U$, $P_j \in \mathscr{P}$, $em_k \in E$.

- A *retrieval* $f(r_x)$ refers to the output obtained when applying the retrieval algorithm to user review r_x. The output is a perfume $P_k \in \mathscr{P}$. The definition of the retrieval is:

$$f\colon TstR \to \mathscr{P}, f(r_x) = P_k \tag{8}$$

- *Total number of relevant retrievals* $TNRR$ refers to the number of outputs which are identified as relevant across all dimensions.

- *Total number of relevant retrievals across dimension* $TNRR_{dim}$ refers to the number of outputs which are identified as relevant across a specific dimension *dim*.

A recommendation is considered relevant if the similarity between the reviewed perfume and the retrieved candidate is above a certain $similarity\_threshold$. Three cases for computing the similarity of two perfumes P_x and P_y are considered:

(1) Their Accords (ACC) are similar. The Accords Similarity (ACCS) criterion refers to a positive result of the comparison between the olfactory profiles of the perfumes, considering that having similar accords, and consequently a similar smell, leads to a match in user preferences.

(2) Their Emotional Vectors (EM) are similar. The Emotional Similarity (EMS) criterion suggests when the emotions identified in the perfumes reviews are alike, i.e. they indicate similar emotional states, atmospheres or moods that the perfumes create to the reviewers.

(3) Their Lifestyle Characteristics (LC) are similar. The Lifestyle Characteristics Similarity (LCS) criterion suggests when the two perfumes are likely to be worn in the same practical and experiential daily activities, thus the two scents would align to the same user lifestyle. The lifestyle characteristics can be divided considering season (SS), time of day wear (TD), perfume longevity (LG), perfume sillage (SL) and perfume gender (GD).

In this study, relevance is operationalized through thresholded similarity across ACCS, EMS and LCS. Importantly, retrieval is driven only by ACCS and EMS, whereas LCS is excluded from ranking and used only as an external validation signal, to prevent circular bias.

The quality of the multi-attribute perfume representation system is evaluated using three complementary metrics that assess different aspects of retrieval performance.

Precision@1 evaluates the quality of the top-ranked retrieval result and is defined as aggregated multi-dimensional similarity

$$Precision@1 = \frac{TNRR}{TstR} \tag{9}$$

Recall@1 evaluates the coverage of the top-ranked retrieval result and is defined as the proportion of relevant recommendations retrieved

$$Recall@1 = \frac{TNRR_{@1}}{TNRR} \tag{10}$$

Traditional precision metrics for similarity retrieval are adapted by defining relevance based on multi-dimensional feature matching. We define the performance measure *Multi-Dimensional Match Rate* ($MDMR$) as the proportion of retrieved candidates that achieve sufficient similarity to the query across individual feature dimensions

$$MDMR = \frac{TNRR_{dim}}{TstR} \tag{11}$$

where dim $\in$ {ACCS, EMS, LCS, SS, TD, LG, SL, GD}

For transparency, a reproducibility and data-availability pipeline is provided in the supplementary material [20].

4.3. Experimental results and their discussion

To validate the model, an ablation study is conducted, comparing the multi-dimensional fusion against two unimodal baselines: ACC-only retrieval ($w = 1$) and EM-only retrieval ($w = 0$). The sensitivity of the fusion model is further analyzed by exploring the impact of the weighting factor w: 0.5 (equal fusion), 0.25 (affective-dominant) and 0.75 (olfactory-dominant).

The primary hypothesis is that the weight factor w significantly affects retrieval process, such that giving slightly more importance to emotions (w=0.25) improves user satisfaction or relevance metrics compared to equal weighting (w=0.5), whereas giving higher importance to accords (w = 0.75) may reduce emotional alignment. Furthermore, an analysis is performed to determine how modifications to the weight parameter w correlate with the system's proficiency in retrievals that align with users' lifestyle characteristics, considering that the perfume discovery is driven only by accord profiles and their corresponding emotional representation.

The empirical results demonstrate that while the $Precision@1$ is highly sensitive to the similarity threshold (Table 7 of the supplementary material [20]), exhibiting a 34% relative decrease as the threshold moves from 0 to 1, $Precision@1$ remains remarkably stable across weighting factor w, suggesting a perfume's emotional profile closely aligns with its olfactory profile.

Similarly, $Recall@1$ exhibited invariance to fluctuations in the weighting factor w, maintaining a consistent value of approximately 0.2. This demonstrates the model's ability to identify a highly relevant match in a single attempt.

Peak performance at $w = 0.25$ indicates that emotional similarity carries the primary predictive signal for fragrance discovery, while a modest olfactory profile (25%) provides complementary structure that improves ranking stability. This result indicates that emotional perception better captures discovery intent than relying on olfactory similarity alone.

The first experiment used $w = 0.5$, i.e. considering equaly importance between accords and emotions similarity. Table 8 of the supplementary material [20] lists average TNRR values over 5 iterations across varying *similarity_threshold*, for the different characteristics: accords (ACCS), emotions (EMS) and lifestyle: season (SS), time of day wear (TD), perfume longevity (LG), perfume sillage (SL), perfume gender (GD) and their combination (ALL).

Interestingly, the emotional vector similarity EMS consistently identified a large number of relevant retrievals across all thresholds. This suggests that emotion similarity plays a significant role in the retrieval process, potentially overshadowing the contributions of the accords similarity.

Increasing the similarity threshold reveals a clear trade-off: higher thresholds result in a smaller number of relevant retrievals (TNRR decreases), but they lead to more focused retrievals.

The significant decrease in TNRR at higher thresholds suggests a potential loss of similarity in retrieval process, i.e. preserving the comfort of familiar scents for the user, thus encouraging the user to explore new scents through peers insights.

In case of ACCS and LCS, quite different optimum threshold is identified, where the balance between TNRR and N-TNRR is optimized. For instance, for ACCS, a threshold around [0.1, 0.2) appears to offer a substantial number of relevant retrievals and maintain a reasonable level of specificity, while with LCS the threhold is around [0.5, 0.6).

At the lowest threshold, all candidates are identified as relevant, in case of each individual LC and the combination of LC. Observe that at rest of thresholds, every individual LC (SS, TD, LG, SL, GD) shows a higher number of relevant retrievals than the their combination (ALL). This indicates that combining all lifestyle characteristics results in a more restrictive filtering with regards to a candidate being identified as relevant.

GD shows slower decrease in the TNRR, which suggests that gender-based perfume preferences are highly specific. LG and SL show a moderate decrease in candidates' relevance with increasing threshold, which suggests that the longevity and sillage characteristics of a perfume do not represent such drastic filter criteria as gender. Similarly, SS and TD yield a high number of relevant retrievals across thresholds, with TD showing a slightly higher decrease as the threshold increases, indicating that time of day wear might be a more specific characteristic than season.

Summing up the observations, evaluating the relevant retrievals based on the concatenation of the lifestyle characteristics proves to be more restrictive in evaluating a retrieval as being relevant, while the individual LC highlight the importance of specific preferences.

Fig. 3 of the supplementary material [20] illustrates how the $MDMR$ measure varies across the weighting factors $w = 0, 0.25, 0.5, 0.75$, respectively 1.

For ACCS, a slight increase of $MDMR$ is observed with the increase of weighting factor w, which is expected because by increasing w, the relative importance of olfactory profile ratio is increasing, leading to retrievals with similar olfactory profiles (Fig. 4 of the supplementary material [20]). For example, for *similarity_threshold* $\in [0.2, 0.3)$, $MDMR$ increases from 0.485 ($w = 0$ and $w = 0.25$) to 0.486 ($w = 0.5$) and 0.487 ($w = 0.75$ and $w = 1$). As a result, 0.3% more retrievals were relevant when increasing w from 0.1 to 0.9 (Fig. 4a).

In case of EMS (Fig. 4b), $MDMR$ remains the same across *similarity_threshold* values, but a slight decrease of relevant retrievals is remarked with increase of weighting factor w, which could be anticipated considering that by increasing w, the relative importance of emotional representation is decreased. $MDMR$ decreases from 0.992 ($w = 0$, $w = 0.25$ and $w = 0.5$) to 0.991 ($w = 0.75$ and $w = 1$). This means that number of relevant retrievals decreases by 0.1%.

When it comes to lifestyle characteristics, it is observed that the impact of weighting factor w systematically enhances the $MDMR$ measure. This pattern is consistent across all lifestyle characteristics Season (Fig. 4d), Time of Day (Fig. 4e), Longevity (Fig. 4f), Sillage (Fig. 4g), Gender (Fig. 4h) and across all similarity thresholds (Fig. 4c). The improvement is particularly noticeable at higher similarity thresholds.

These results underline the importance of balancing olfactory profile and emotional characteristics when modeling perfume similarity. The consistent increase in $MDMR$ across all thresholds with increase of w shows that accentuating olfactory profile improves the system's ability to retrieve lifestyle-appropriate perfume discovery (accords encode richer information for lifestyle matching, while emotional descriptors contribute complementary nuance).

5. Conclusions

This paper proposed a multi-dimensional model for fragrance discovery, using fragrance profile and emotions extracted from user reviews. The system receives as input a perfume review and retrieves relevant perfume candidates according to user preferences for the reviewed perfume. A weighted aggregation is employed to balance the influence of olfactory and affective similarity scores, controlled by weighting parameter w.

To evaluate the retrieval quality, evaluation measure *Multi-Dimensional Match Rate* ($MDMR$) is proposed. $MDMR$ shows the part of retrievals which are identified as relevant across three distinct facets: olfactory (ACCS), affective (EMS), and contextual (LCS). Crucially, Lifestyle

Characteristics (LC) were withheld from the retrieval algorithm and used as an external dimension to assess whether accord-emotion-based retrievals align with contextual suitability.

It is observed that olfactory alignment (ACCS) remains robust even at lower similarity thresholds, allowing for flexible discovery without losing scent consistency. Furthermore, the high degree of affective alignment (EMS) across all weight factors confirms that social-emotional signals provide a stable anchor for fragrance discovery. However, the current reliance on a dominant 6-dimensional emotional model (Parrott) may lead to high profile similarity; therefore, investigating more granular affective models to enhance clustering remains a priority for future work.

The evaluation against Lifestyle Characteristics (LCS) provided a crucial external validation of the system's performance. While individual facets (such as seasonality, longevity, or sillage) show high baseline alignment, the concatenated lifestyle profile (ALL) serves as a significantly more rigorous and restrictive filter. The high number of relevant retrievals maintained across this multi-dimensional constraint demonstrates that the system does not merely match surface-level metadata. Instead, fusing olfactory and affective data captures a fragrance's utility, predicting environmental suitability even when lifestyle attributes are withheld from the ranking logic.

References

[1] J. M. GABY and P. DALTON, *Discrimination between individual body odors is unaffected by perfume*, Perception **48**(11), 2019, pp. 1104–1123.

[2] L. MARCOTULLIO, P. GROB, S. SCIALLA and P. SAVEYN, *Understanding perfume deposition mechanisms on different textile substrates in a realistic laundering process*, Journal of Surfactants and Detergents **28**(4), 2025, pp. 731–745.

[3] J. LIU, R. SUN, X. BAO, J. YANG, Y. CHEN, B. TANG and Z. LIU, *Machine learning driven atom-thin materials for fragrance sensing*, Small **21**(28), 2025, paper 2401066.

[4] E. UPSTILL, C. ROWLAND and C. JORDAN, *Delivering context for fragrance evaluation: A study using trained sensory panellists*, Food Quality and Preference **102**, 2022, paper 104636.

[5] J. MENSING, *Psychology of perfume choice*, in Beautiful SCENT: The Magical Effect of Perfume on Well-Being, Springer Berlin Heidelberg, Berlin, Heidelberg, 2023, pp. 53–80.

[6] J. MENSING, *Perfume trends, international fragrance preferences and mentalities*, in Beautiful SCENT: The Magical Effect of Perfume on Well-Being, Springer Berlin Heidelberg, Berlin, Heidelberg, 2023, pp. 319–335.

[7] R. S. HERZ, *Chapter 17: Perfume*, in Neurobiology of Sensation and Reward, J. A. Gottfried, Ed., CRC Press, 2011, pp. 371–390.

[8] E.-R. LUTAN, *Challenges in perfume recommender systems: navigating subjectivity, context and sensory data*, Proceedings of the 19th ACM Conference on Recommender Systems (RecSys '25), Prague, Czech Republic, 2025, pp. 1469–1472.

[9] D. M. V. ZEMKE and S. SHOEMAKER, *Scent across a crowded room: Exploring the effect of ambient scent on social interactions*, International Journal of Hospitality Management **26**(4), 2007, pp. 927–940.

[10] N. KHAIR, S. ELHAJJAR and Z. HAMZEH, *Personal branding through perfume in the Middle East: Investigating the role of fragrance in self-presentation, impression management, and cultural identity*, Fashion Theory - The Journal of Dress Body & Culture **28**(4), 2024, pp. 461–481.

[11] Y. ZHONG, Y. LI, H. ZHOU and L. ZOU, *What Perfume Makes Us Feel: A self-report evaluation of perfume-related olfactory aesthetic experience scale*, Psychology of Aesthetics, Creativity, and the Arts **20**(3), 2024, pp. 599–608.

[12] K. SU, C. WANG, L. HU, X. DAI, X. FANG, C. YANG, R. ZHANG, Q. LI, R. ZHOU, Z. LU, D. LI, Y. XIE, C. ZHOU, Z. ZHAO, X. CHEN, H. WANG and J. XIANG, *Indoor use of consumer chemical products and its link to depression and anxiety among college students in 297 Chinese cities*, Building and Environment **279**, 2025, paper 113057.

[13] M. SJARMA, S. BALAJI, P. SAHA and R. KUMAR, *Navigating the fragrance space using graph generative models and predicting odors*, Journal of Chemical Information and Modeling **65**(10), 2025, pp. 4818–4832.

[14] P. HANAFIZADEH, A. ZARERAVASAN and H. R. KHAKI, *An expert system for perfume selection using artificial neural network*, Expert Systems with Applications **37**(12), 2010, pp. 8879–8887.

[15] O. PIMENTEL, E. G. CHUQUICHAMBI, C. SPENCE and C. VELASCO, *The diatonic sound of scent imagery*, Perception **54**(9), 2025, pp. 689–714.

[16] F. A. T. TOBING, A. KUSNADI, E. E. SURBAKTI, W. R. HARSONO, A. I. WIJAYA, F. DINARTA, M. IRELYNN and R. PASKAH, *Men's perfume recommendation system using analytic hierarchy process (AHP) and technique for order of preference by similarity to ideal solution (TOPSIS) method*, International Journal of Science Technology & Management **4**(6), 2023, pp. 1552–1559.

[17] C. PORCHEROT, S. DELPLANQUE, S. RAVIOT-DERRIEN, B. LE CALVE, C. CHREA, N. GAUDREAU and I. CAYEUX, *How do you feel when you smell this? Optimization of a verbal measurement of odor-elicited emotions*, Food Quality and Preference **21**(8), 2010, pp. 938–947.

[18] S. R. JAEGER, S. M. LEE, K.-O. KIM, S. L. CHHEANG, C. M. ROIGARD and G. ARES, *CATA and RATA questions for product-focused emotion research: Five case studies using emoji questionnaires*, Food Quality and Preference **68**, 2018, pp. 342–348.

[19] H. TOUMY, *Perfume Project*, 2019. Accessed: Jan. 20, 2026. [Online]. Available: https://hayatoumy.github.io.

[20] E.-R. LUȚAN and C. BĂDICĂ, *Supplementery material of the paper Modeling Fragrance Discovery through Multi-Attribute Fusion of Olfactory Accords and Affective Profiles*, 2026. [Online] https://drive.google.com/file/d/1Icn0KM3I74otTxz1vbNoiKoSyp5nAWEQ/view?usp=sharing.

[21] E.-R. LUȚAN, C. BĂDICĂ and N. I. ENESCU, *Cross-domain emotion-based recommender system for books and movies*, Proceedings of 2024 International Conference on INnovations in Intelligent SysTems and Applications (INISTA), Craiova, Romania, 2024, pp. 1–6.

[22] E. SARAVIA, H.-C. T. LIU, Y.-H. HUANG, J. WU and Y.-S. CHEN, *CARER: Contextualized affect representations for emotion recognition*, Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, 2018, pp. 3687–3697.