Received: 28 July 2023 | Revised: 15 November 2023 | Accepted: 30 November 2023

DOI: 10.1002/hup.2889

RESEARCH ARTICLE

WILEY

Trajectories of sentiment in 11,816 psychoactive narratives

Sam Freesun Friedman[1] | Galen Ballentine[2]

[1]Data Sciences Platform, Broad Institute of MIT and Harvard, Cambridge, Massachusetts, USA

[2]Department of Psychiatry, SUNY Downstate Health Sciences University, Brooklyn, New York, USA

Correspondence
Sam Freesun Friedman.
Email: sam@broadinstitute.org

Abstract

Objective: Can machine learning (ML) enable data-driven discovery of how changes in sentiment correlate with different psychoactive experiences? We investigate by training models directly on text testimonials from a diverse 52-drug pharmacopeia.

Methods: Using large language models (i.e. BERT) and 11,816 publicly-available testimonials, we predicted 28-dimensions of sentiment across each narrative, and then validated these predictions with adjudication by a clinical psychiatrist. BERT was then fine-tuned to predict biochemical and demographic information from these narratives. Lastly, canonical correlation analysis linked the drugs' receptor affinities with word usage, revealing 11 statistically-significant latent receptor-experience factors, each mapped to a 3D cortical Atlas.

Results: These methods elucidate a neurobiologically-informed, sequence-sensitive portrait of drug-induced subjective experiences. The models' results converged, revealing a pervasive distinction between the universal psychedelic heights of feeling in contrast to the grim, mundane, and personal experiences of addiction and mental illness. Notably, MDMA was linked to “Love”, DMT and 5-MeO-DMT to “Mystical Experiences” and “Entities and Beings”, and other tryptamines to “Surprise”, “Curiosity” and “Realization".

Conclusions: ML methods can create unified and robust quantifications of subjective experiences with many different psychoactive substances and timescales. The representations learned are evocative and mutually confirmatory, indicating great potential for ML in characterizing psychoactivity.

KEYWORDS
large language models, machine learning, psychoactives, subjective experiences, testimonials

1 | INTRODUCTION

Psychoactive molecules can engender an awesome breadth of subject states—from inebriated to hallucinatory, mystical to numb—with massive pharmaceutical and societal implications (Shulgin & Shulgin, 1991). Critically, however, researchers lack a unified framework for quantifying how subjective drug effects unfold. The instruments used to measure the subjective effects of psychoactive drugs consist primarily of questionnaires (e.g. DEQ, 5D-ASC, ARCI) (Haertzen, 1966; Hasler et al., 2004; Kubany et al., 2000) or symptom scales (e.g. HAMD, YBOCS) (Moritz et al., 2002; Williams et al., 2008). The emergence of artificial neural networks alongside growing datasets of drug testimonials provides an opportunity to develop powerful new methods for quantifying psychoactivity.

Eschewing methodological boundaries between psychiatric medications, substances of abuse, and hallucinogens, we pool together qualitative data revealing patterns of drug induced experience across drug classes. This exploratory, descriptive study presents an AI-informed, data-driven perspective on how different drugs associate with different sentiments and semantic tags as well as the trajectories these associations follow over the course of narrative testimonials.

Three distinct but related machine learning (ML) modeling approaches are brought to bear on a corpus of 11,816 drug narratives from Erowid. Firstly, a large language model was trained on publicly available, human-annotated texts to predict 28-dimensions of sentiment which were then validated with adjudication by a clinical psychiatrist. Secondly, a separate model was fine-tuned to predict biochemical and demographic information from these narratives, and then generalized to unseen data. Thirdly, canonical correlation analysis (CCA) linked the drugs' receptor affinities with word usage, revealing 11 statistically-significant latent receptor-experience factors, each mapped to a 3D cortical Atlas. Together, these three ML methods elucidate a neurobiologically-informed, sequence-sensitive portrait of drug-induced subjective experiences.

Gross aspects like arousal and hedonic tone are conserved across our species, while diverse subjective manifestations of psychiatric illnesses, intensification of emotion during acute drug (Nichols, 2016) experiences, and addiction require more nuanced, multi-dimensional classifications (Russell & Barrett, 1999). Recent advances in Natural Language Processing (NLP) provide new means to quantify this nuance by manually-annotating taxonomies of emotion in large corpora of text (Dang et al., 2020; Demszky et al., 2020; Maas et al., 2011). Simultaneously, transformer-based models have made tremendous progress on many language tasks (Brown et al., 2020; Devlin et al., 2018; Vaswani et al., 2017).

Feelings unfold over time. This temporal trajectory of sentiment is the substrate in which significance is felt. Moreover, the form of a sentiment's trajectory shapes the memory that is ultimately consolidated. For example, the "peak-end rule" states that the intensities at the peak and the end of an experience determine its imprint in memory, while the notion of "primacy" holds that earlier events in a sequence are remembered more clearly (Altmann, 2000; Kahneman et al., 1993). Moreover, emotion can modulate the effects of primacy and recency (Forgas, 2011). Such subtleties become essential when considering the therapeutic effect of psychedelic-assisted therapy wherein the observed correlation between mystical experience and clinical efficacy is likely to be conditioned by the precise sequential dynamics and emotional content of the acute experience (Yaden & Griffiths, 2021).

Traditional questionnaires are not designed to capture such trajectories. While questionnaires are indispensable, validated instruments for human research, they inevitably compress a vast range of psychoactivity to selected aspects of an experience (e.g. Five items for the DEQ, 30 for the Mystical Experience Questionnaire (F. S. Barrett et al., 2015), 85 for the Hallucinogen Rating Scale, 94 for the 5D-ASC, 550 for the ARCI). The complex dynamics of feelings throughout a drug experience are reduced to low-dimensional points in a space spanned by the researcher's conceptual priorities, rather than those of the subject. In contrast, subject testimonials paired with large language models can rigorously quantify many dimensions of subjective experience as they evolve throughout a narrative report. Here, we demonstrate three distinct modeling techniques (i.e. self-supervised, transferred and inferred) that all discover overlapping and neurochemically-grounded structures from subjective testimonials of drug experiences. These structures richly depict retrospective associations, but are not necessarily predictive or mechanistic.

Our contributions with respect to prior work (Coyle et al., 2012; Sanz et al., 2018; Zamberlan et al., 2018) include: (1) extending CCA (Ballentine et al., 2022) to 52 drugs and 11,816 narratives, eliminating affinities whose replicability had been questioned (Galloway, 2022) by exclusively sourcing from the Psychoactive Drug Screening Program (PDSP) (Roth et al., 2000) and thus registering a cacophony of pharmaceuticals into a unified, neurobiologically-grounded space of psychoactivity; (2) leveraging pretrained language models and transfer-learning we construct trajectories, rather than static summaries, discerning subtle similarities and differences between individual drugs and drug classes; (3) predicting demographics directly from natural language demonstrating the feasibility of automated bias-reduction methods; and (4) showing that these diverse modeling strategies all independently elucidate a striking structure that juxtaposes the mystical, beautiful, universal heights of psychedelia against the grim, personal and painful battles with addiction and mental illness.

2 | METHODS

Many datasets were leveraged in this study, namely Erowid (EROWID: Documenting the Complex Relationship Between Humans & Psychoactives, 2000), receptor affinities at 61 receptor subtypes for 44 drugs from the PDSP(Besnard et al., 2012; Roth et al., 2000) (Supplementary Figure 10), augmented with affinities for 8 phenethylamines from Rickli et al., 2015), RNA gene expression data for 200 brain regions from the Allen Brain Atlas (Sunkin et al., 2013), and 58K Reddit posts with 28 human-annotated human emotions (Demszky et al., 2020), see Supplementary Figure 20 for a schema. The 10 pharmacologic classes and the 22 chemical classes were retrieved from the Psychonaut Wiki (Schifano et al., 2006). The code necessary to run these analyses are available here: https://github.com/lucidtronix/bertowid.

Testimonials were downloaded from Erowid (www.erowid.org), an educational resource on the effects of psychoactive substances—both legal and illegal. The collection includes >42,000 reports in English from >800 different drugs. Testimonials are submitted anonymously and without financial incentive. Quality control is conducted by trained experts with careful consideration given to "quality, credibility, and focus on effects or outcomes" (EROWID:

Documenting the Complex Relationship Between Humans & Psychoactives, 2000). Many testimonials are recounted chronologically with sequential timestamps, while others include temporally ambiguous information like drug preparation details, social context and philosophical waxings. The quality controls of Erowid and the well-established genre of trip reports dating back to Aldous Huxley's *Doors of Perception* (Huxley, 2014) and Carlos Castañeda's *Teachings of Don Juan* (Castaneda, 1998), ensures that, though they do not map directly to time, the testimonials are consistently structured narrative sequences. These narratives therefore contain information beyond the precise chronology of the original experience.

Pharmacologic and chemical taxonomies provide the biochemical anchor for the phenomenology described in the testimonials. Pairing affinity data at each receptor subtype with RNA expression levels in the cortex allows us to localize the semantic structures to specific brain regions. So, using only publicly-available data and pretrained large language models we show it is possible to build neurobiologically-informed sequential trajectories of drug-induced subjective experiences.

2.1 | Data preprocessing for transformers

Scraped testimonials were parsed for meta data, drug-masked, and tokenized. Drug-masking removed all occurrences of drug names in the testimonial text, including both scientific, common, and colloquial nomenclature as well as misspellings. See Supplementary Table 3 and code for the full list of masked words. Models were initialized with pretrained weights for the base BERT encoders. All initial model weights are publicly available. Except when otherwise noted, the base BERT model used was trained with the Stanford Sentiment Treebank data, a corpus of natural language movie reviews with each review annotated for sentiment valence and degree by humans (Socher et al., 2013). The initial model architecture and weights are available at TensorFlow Hub. The pooled output from the base BERT model was extended with a dropout layer (Srivastava et al., 2014) followed by a dense layer for each task (e.g. BERTiment has 28 distinct outputs–one for each binary emotion classification: present or absent).

While testimonials vary in size, the input to BERT models is at most 512 tokens. A sliding window inference step used all available data by creating prediction series of varying lengths for each testimonial from each model. Different window sizes are compared in Supplementary Figure 6. When the window size exceeds the testimonial size, the input is zero-padded. When the window size is smaller than the testimonial size, the testimonial is split into contiguous blocks of text on newlines and sentence ending punctuation, and the model is applied to each. In this way, a trajectory of inferences is made. Dynamic Time Warping (DTW) quantified inter-trajectory distances, using an implementation from the fastdtw python package (Salvador & Chan, 2007). The Broad Institute's ML4H tools were used for model evaluation and tensor-mapping (Friedman et al., 2020; Sarma et al., 2020).

2.2 | BERT-based model encoder fine-tuning

BERT stands for Bidirectional Encoder Representations from Transformers, a type of large language model which learns an embedding (i.e. representation) for text which can then be applied to many different language modeling tasks, such as emotion classification and demographic prediction. The encoder backbone of both BERTowid and BERTiment was trained with a masked language model objective (Devlin et al., 2018; Taylor, 1953). BERTowid and BERTiment add output heads, which take the BERT encoder representation as input and fine-tune it for new tasks. The encoder backbone contains 109 million parameters. Base models are compared in Supplementary Figure 7, showing similar performance with different pre-training datasets. Prior to the output layer for the fine-tuning task, we insert a dropout layer, see Supplementary Figure 8 (Srivastava et al., 2014). The ADAMw (Loshchilov & Hutter, 2017) stochastic gradient descent optimization with initial learning rate of 1e-5 and a batch size of 32 is used for fine-tuning. The minimum validation loss model is serialized for downstream inference after 16 epochs.

2.3 | Training BERTowid

The Erowid metadata, the inferred CCA weights and the receptor affinities from PDSP provide diverse training labels for BERTowid. Both classification (e.g. drug, tag, gender) and regression tasks (e.g. age, affinity, CCA weights) are considered. All classification tasks are trained to minimize a cross entropy loss, while regression models are trained to minimize the mean squared error of their predictions. Classification and regression with a single model requires a term to balance between the two types of loss, but optimization was found to be sensitive to this value, requiring careful tuning to avoid convergence for only one of the loss types. To mitigate this, multitask BERTowid is trained and serialized separately for classification and regression. Supplementary Figure 16 compares multi-task versus single task models showing a relatively small cost to taking the multitask approach. Supplementary Figure 17 shows results with weighted loss functions to mitigate class-imbalance. Less common tags are in fact less informative, perhaps because they are less rigorously or consistently attached by the Erowid moderators. Giving these less informative labels more weight in the loss function results in a worse model. With a window size of 64 words and all categorical tasks, 16 epochs takes about 3 h on a NVidia V-100 GPU.

2.4 | Training BERTiment

BERTiment is trained to predict the 28 emotion classifications which are annotated in 58K Reddit posts in the Google Emotions data set, which can be downloaded here. The BERTiment training procedure has previously been described and evaluated in the GoEmotions paper (Demszky et al., 2020). Our approach only differs in dropout rate, 0.5, learning rate, 1e-5, and batch size, 32,

where for consistency we used the same hyper-parameters and base model used to train BERTowid. We split the data 70-20-10 between training, testing and validation sets. The 28-head BERTiment is trained for 16 epochs which takes about 2 h on a NVidia V-100 GPU.

2.5 | Canonical correlation analysis

The CCA mapping between testimonials and affinities has previously been described in detail (Ballentine et al., 2022). The approach here extends from 27 drugs to 52, from 40 receptor subtypes to 61, and exclusively sources affinity values from the PDSP(Besnard et al., 2012) or Rickli et al., 2015). Briefly, CCA is applied to two linked matrices, one contains natural language data, the other receptor affinity data. Both matrices have a row for each testimonial. After under-sampling from over-represented drugs, 11,141 testimonials are included in the CCA analysis. The natural language matrix is made from PCA-reduction on bag-of-words representations of the testimonials, specifically from the tf-idf transformed word count matrix of 20,943 words, the top 1000 principal components are kept. So the natural language matrix has dimensions 11,141 by 1000. The receptor affinity matrix has the 61 affinity values for the drug tiled to link to each testimonial and so has dimensions 11,141 by 61. CCA reveals patterns of correlation that link word usage to receptor binding strengths. Each of these receptor-semantic components is composed of two poles of a weighted list of words and a weighted list of receptors. The receptor weights for each component were mapped to the cortex using Allen Brain Atlas receptor RNA expression quantities measured by invasive tissue probes. The scikit-learn package's implementation of the CCA algorithm was used (Pedregosa et al., 2011).

For each of the 52 drugs to be analyzed, we built an affinity vector that captures the binding strengths (Ki) for 61 targets including G protein–coupled receptors, molecular transporters, and ion channels (Supplementary Table 4). Binding assays performed in coordination with the PDSP followed the methodology of Glennon et al. (53): Briefly, for each compound, a primary assay at 10 nM concentration was performed against each receptor, transporter, or ion channel. Those compounds that induced a "hit" of >50% inhibition were then subjected to a secondary assay at 1, 10, 100, 1000, and 10,000 nM to determine Ki values, with the final value calculated as the average of at least three replicated assays. Further details of how individual assays were conducted can be found at https://pdsp.unc.edu/databases/kidb.php.

We lastly anatomically locate the significant receptor-experience factors based on their gene transcription weighting of neurotransmitter receptors. Publicly available human gene expression data from six whole postmortem brains of neurotypical donors were obtained from the Allen Human Brain Atlas (http://human.brain-map.org). To strengthen reproducibility and comparability, we have used the abagen tool to map the receptor gene expression information to the 200 Schaefer-Yeo regions (https://github.com/rmarkello/abagen). Averages of invasive brain tissue probes were computed across all six donors for each of the 61 receptors of interest. The factor-specific coexpression of receptor genes was topographically mapped to each of the 200 target brain regions of the Schaefer-Yeo reference Atlas.

3 | RESULTS

We amassed a corpus of 11,816 psychoactive experiences, which we semantically and chemically characterize with two BERT-based models, and one Canonical Correlation Analysis (CCA) see Figure 1. The transfer-learning model, BERTowid, is trained using multi-task, multi-label, classification and regression directly on Erowid testimonials and associated metadata. BERTowid is trained to "read" a 512 token excerpt from the testimonial and predict the associated drug, its chemical and pharmacological class, self-reported gender and age, 52 metadata tags, 11 canonical correlation component weightings, and 30 receptor affinities. Table 1 shows taxonomy information for each drug and Table 2 shows the testimonial counts. A second model, BERTiment, is pretrained on a corpus from Reddit to detect 28 sentiments simultaneously (Demszky et al., 2020). Inference on Erowid then reveals sentimental trajectories which we validate with a psychiatrist's adjudications. The predicted sentiment distributions also confirm emotional associations recognized by drug subcultures and are consistent with pharmacological groupings. Both models generalize to unseen data, demonstrating how machine learning on crowd-sourced, noisy data can lead to diverse biochemical inferences. Note for instance in Figure 2 how the entactogens MDA and MDMA, the opioids and the antidepressants all track together. The mean trajectories for a given drug reliably and reproducibly segregate drug experiences and conform to cultural expectations of different substances. Simultaneously, the shaded regions surrounding the mean trajectories indicate +/− a standard deviation in predicted sentiment within each drug. The large overlap in shaded regions speaks to variation of sentiments described even for a given drug. This is no surprise as it has long been recognized that set and setting can dramatically impact the subjective experience with the same psychoactive molecule. Comparison of the mean trajectories between drugs gives a population-level inter-drug characterization of a substance, while trajectories from a single testimonial can characterize individuals by showing how their trajectory stacks up against the population mean. The most extreme individual testimonial for each sentiment and semantic tag are shown in Supplementary Tables 1 and 2.

Another population-level characterization is provided by CCA, which identified a latent structure of 11 statistically-significant components mapping between the semantic data and the receptor affinity profiles in a self-supervised fashion (Ballentine et al., 2022). CCA is a linear model and relies on a bag-of-words representation of the entirety of the testimonial text, while the transformers are deep nonlinear neural networks which positionally-encode a subset of text

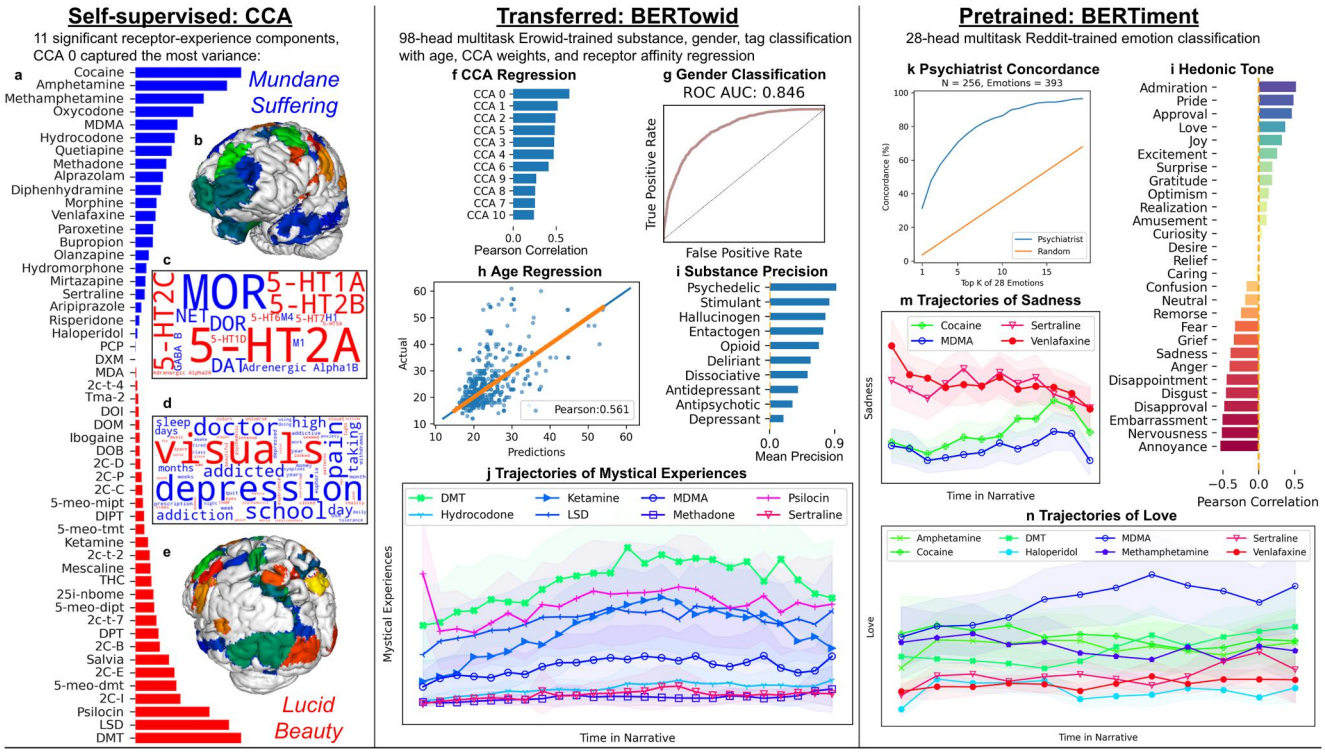

FIGURE 1 Models. This figure illustrates the three models with a selection of key results. Left: the dominant component, CCA 0, found by the self-supervised learning CCA method. One extreme of CCA 0 encodes concepts of somatic suffering displayed in blue, while the other pole encompasses visual beauty and is displayed in red. All 52 drugs included in the study are shown in (a) with the ranking along CCA 0. (b) Shows a brain surface map of CCA 0. (c) Shows receptor clouds, note how the red visual/beauty pole is entirely serotonergic (5HT*) while the somatic/suffering pole in blue highlights several different neurotransmitter types including opioid (MOR, DOR), GABA and acetylcholine (M1, M4). (d) Shows word clouds with font size determined by their CCA 0 weighting. (e) Shows a surface view of CCA 0 mapping into the brain, note how the visual pole highlights the visual cortex. Middle: results from the transfer learning model, BERTowid, which is a multi-task classification and regressing transformer trained directly on Erowid testimonials. All results are from test set testimonials, which were not used in training. (f) Pearson correlation with the 11 CCA factors per-testimonial weightings. (g) Self-reported gender ROC curve. (h) Pearson correlation with self-reported age. (i) Mean precision per psychoactive class. Tiling inferences from BERT models along the narrative of the testimonials we construct trajectories, for clarity we only show a few of 52 drugs here. (j) Trajectories for the semantic tag of "Mystical Experiences", note the prominence of DMT. Right: pretrained model results from BERTiment, trained on an entirely different text corpus to classify emotions. (k) BERTiment's concordance with a clinical-psychiatrist emotion adjudication in Erowid testimonials. (l) IMDB movie review hedonic-tone classifier correlation with the 28 emotions inferred on Erowid. (m) BERTiment Sadness trajectories, note how the antidepressants track with each other and are initially quite elevated. (n) BERTiment Love trajectories, note the prominence of MDMA.

excerpted from the testimonials. Despite the large differences in representation and model, BERTowid learns to infer the CCA weightings, while many BERTiment emotion-scapes reveal similar drug rankings as given by CCA 0, see Figure 3.

3.1 | BERTowid

There is noise inherent in any crowd-sourced, open dataset like Erowid, which includes reports from many illegal substances rife with potential impurities and misrepresentations. Nonetheless, BERTowid shows powerful discrimination at several different granularities of pharmacology, classifying among 52 drugs, 22 ligand chemical types, and 10 pharmacologic classes, 30 receptor subtypes, and 11 CCA weights. Model mistakes are consistent with expected pharmacological groupings, for example, the psychedelic chemical classes of phenethylamines and tryptamines are more often mistaken for each other, than for less similar chemicals (see confusion matrices in Supplementary Figure 1). Erowid-supplied 52 semantic tags as metadata, many of which are also learnable, with some of the best-performing being "Medical Use", "Mystical Experiences", "Alone" and "Addiction Habituation", with areas under the receiver operating characteristic curves (ROC AUC) ranging from 0.88 to 0.95, metrics for all tags are shown in Supplementary Figure 2 and Supplementary Figure 18.

Conforming to its reputation as the "spirit molecule", DMT displayed heightened levels of "Mystical Experiences" and, even more dramatically, for the tag "Entities and Beings", echoing themes uncovered in manual DMT-specific analyses (Lawrence et al., 2022). As expected, the "Depression" tag highlights antidepressants, while the "Addiction Habituation" tag is consistently elevated for the stimulants cocaine and methamphetamine, see Supplementary Figure 3.

TABLE 1 Drug overview.

Pharmacological class	Drug name	Common and brand names	Chemical class	Non-Medical dosage range (in milligrams)	Therapeutic dosage range (in milligrams)	Average duration of action (in hours)	Route(s) of administration	SMILES	Sources
Stimulant	Amphetamine	*Speed, Adderall, Pep*	Phenethylamine	2.5–50	5–40	5	Oral, insufflation	CC(CC1=CC=CC=C1)N	Pharmacologic class, common/Brand names, chemical class, routes of administration, non-medical dosage range, average duration of action values were taken rom psychonautwiki.org; herapeutic dosage range values were taken from UptoDate.com; SMILES were taken from https://pubchem.ncbi.nlm.nih.gov/
	Bupropion	*Wellbutrin, Zyban, Aplenzin*	Cathinone	75–325	100–450	10	Oral	CC(C(=O)C1=CC(=CC=C1)Cl)NC(C)(C)C	
	Cocaine	*Coke, Coca, Crack, Blow, Girl, White, Snow, Nose Candy*	Tropane alkaloid	5–90	-	0.82	Insufflation, IV, oral	CN1C2CCC1C(C(C2)OC(=O)C3=CC=CC=C3)C (=O)OC	
	Methamphetamine	*Meth, Speed, Ice, Glass, Shard, Tina, Crank, Desoxyn*	Amphetamine	5–60	20–60	9.5	Oral, insufflation, IV	CC(CC1=CC=CC=C1)NC	
Psychedelic	25i-NBOME	-	Phenethylamine	1–16	-	15	Oral	COC1=CC=CC=C1CNCCC2=CC(=C(C=C2OC)I)OC	
	2C-B	*Nexus, Bees*	Phenethylamine	5–45	-	6.5	Oral	COC1=CC=CC=C1CNCCC2=CC(=C(C=C2OC)I)OC	
	2C-C	-	Phenethylamine	5–90	-	6	Oral	COC1=CC(=C(C=C1CCN)OC)Cl	
	2C-D	*2C-M, LE-25*	Phenethylamine	3-00	-	4	Oral	CC1=CC(=C(C=C1OC)CCN)OC	
	2C-E	*Eternity, Aquarust*	Phenethylamine	2–30	-	8	Oral	CCC1=CC(=C(C=C1OC)CCN)OC	
	2C-I	-	Phenethylamine	2–30	-	8	Oral	COC1=CC(=C(C=C1CCN)OC)I	
	2C-P	-	Phenethylamine	3–30	-	8	Oral	CCCC1=CC(=C(C=C1OC)CCN)OC	
	2C-T-2	*Rosy*	Phenethylamine	3–30	-	8	Oral	CCSC1=C(C=C(C(=C1)OC)CCN)OC	
	2C-T-4	-	Phenethylamine	8–20	-	15	Oral	CC(C)SC1=C(C=C(C(=C1)OC)CCN)OC	
	2C-T-7	*Beautiful, Blue Mystic, 7th Heaven*	Phenethylamine	3–40	-	8	Oral	CCCSC1=C(C=C(C(=C1)OC)CCN)OC	
	5-MeO-DiPT	*Foxy Methoxy, Foxy*	Tryptamine	3–20	-	6	Oral	CC(C)N(CCC1=CNC2=C1C=C(C=C2)OC)C(C)C	
	5-MeO-DMT	*The God Molecule, Toad, Jaguar, Soma*	Tryptamine	1–20	-	0.5	Oral	CN(C)CCC1=CNC2=C1C=C(C=C2)OC	
	5-MeO-MiPT	*Moxy*	Tryptamine	3–20	-	6.5	Oral	CC(C)N(C)CCC1=CNC2=C1C=C(C=C2)OC	
	5-MeO-TMT	-	Tryptamine	75–150	-	7.5	Oral	CC1=C(C2=C(N1)C=CC(=C2)OC)CCN(C)C	
	DiPT	-	Tryptamine	15–150	-	4.5	Oral	CC(C)N(CCC1=CNC2=CC=CC=C21)C(C)C	
	DMT	*N,N-DMT, Dmitry, The Glory, The Spirit Molecule*	Tryptamine	2–60	-	0.2	Inhalation	CN(C)CCC1=CNC2=CC=CC=C21	
	DOB	*Bromamfetamine, Bromo-DMA*	Amphetamine	0.2–3	-	19	Oral	CC(CC1=CC(=C(C=C1OC)Br)OC)N	
	DOI	-	Amphetamine	0.5–3	-	20	Oral	CC(CC1=CC(=C(C=C1OC)I)OC)N	
	DOM	*STP (Serenity, Tranquility, and Peace)*	Amphetamine	0.5–10	-	14	Oral	CC1=CC(=C(C=C1OC)CC(C)N)OC	
	DPT	*The Light*	Tryptamine	50–350	-	3	Oral	CCCN(CCC)CCC1=CNC2=CC=CC=C21	
	Ibogaine	*Endabuse, Iboga*	Tryptamine	15	-	45	Oral	CCC1CC2CC3C1N(C2)CCC4=C3NC5=C4C =C(C=C5)OC	
	LSD	*LSD-25, Lucy, L, Acid, Cid, Tabs, Blotter*	Lysergamide	0.015–0.3	-	10	Oral	CCN(CC)C (=O)C1CN(C2CC3=CNC4=CC=CC (=C34)C2=C1)C	
	Mescaline	*Mescaline, Peyote, San Pedro, Cactus, Buttons, Devil'*	Phenethylamine	50–800	-	11	Oral	COC1=CC(=CC(=C1OC)OC)CCN	

TABLE 1 (Continued)

Pharmacological class	Drug name	Common and brand names	Chemical class	Non-Medical dosage range (in milligrams)	Therapeutic dosage range (in milligrams)	Average duration of action (in hours)	Route(s) of administration	SMILES	Sources
	Psilocin	*Psilocin, Psilocine, Psilocyn, Psilotsin*	Tryptamine	5–40	20–30	5	Oral	CN(C)CCC1=CNC2=C1C(=CC=C2)O	
	TMA-2	-	Amphetamine	5–60	-	10	Oral	CC(CC1=CC(=C(C=C1OC)OC)OC)N	
Opioid	Hydrocodone	*Vicodin, Zohyrdo ER*	Morphinan	3–40	10–80	6	Oral, insufflation, IV	CN1CCC23C4C1CC5=C2C(=C(C=C5)OC)OC3C (C=C4)O[Si](C)(C)C	
	Hydromorphone	*Dilaudid*	Morphinan	0.5–8	1–12	3.5	Oral, insufflation, IV	CN1CCC23C4C1CC5=C2C(=C(C=C5)O) OC3C(=O)CC4	
	Methadone	*Dolophine*	Diphenylpropylamine	1–30	10.0–120	6	Oral	CCC(=O)C(CC(C)N(C)C) (C1=CC=CC=C1) C2=CC=CC=C2	
	Morphine	*MSContin, Oramorph, Zomorph, Sevredol*	Morphinan	10–30	1–10	5	Oral, insufflation, IV	CN1CCC23C4C1CC5=C2C(=C(C=C5)O) OC3C(C=C4)O	
	Oxycodone	*OxyContin, Oxy, Roxicodone, Oxecta, OxyIR, Endone, Oxynor, Codilek, Oxydore, Redocam, Oxygesic*	Morphinan	1–40	2.5–40	5	Oral, insufflation, IV	CN1CCC23C4C(=O)CCC2(C1CC5=C3C (=C(C=C5)OC)O4)O	
Hallucinogen	Salvinorin A	*Salvia, Salvia Divinorum, Diviner's Sage, Ska Maria Pastora, Seer's Sage, Sally*	Salvinorin	0.2–1.0	-	0.9	Smoked	CC(=O)OC1CC(C2(CCC3C(=O)OC(CC3(C2C1 =O)C)C4=COC=C4)C)C (=O)OC	
	THC	*Cannabis, Marijuana, Weed, Pot, Mary Jane*	Cannabinoid	0.4–10	2–20	3.65	Smoked, oral	CCCCCC1=CC(=C2C3C=C(CCC3C(OC2=C1)(C)C)C)O	
Entactogen	MDA	*Sass, Sassafras, Tenamfetamine*	Amphetamine	20–145	-	6.5	Oral	CC(CC1=CC2=C(C=C1)OCO2)N	
	MDMA	*Molly, Mandy, Emma, MD, Ecstasy, E, X, XTC, Rolls, Beans*	Amphetamine	10.0–180	80–120	4.5	Oral	CC(CC1=CC2=C(C=C1)OCO2)NC	
Dissociative	DXM	*DXM, DMO, DM, Dex, Robitussin, Delsym, DexAlone, Duract*	Morphinan	10.0–700	10–60	10	Oral	CN1CCC23CCCCC2C1CC4=C3C=C(C=C4)OC	
	Ketamine	*Ketamine, K, Ket, Kitty, Special K, Cat Tranquilizer, Ketaset, Ketalar, Ketanest, Vitamin K, Purple, Jet*	Arylcyclohexylamine	5.0–1000	25–100	1.25	IV, IM, buccal	CNC1(CCCCC1=O)C2=CC=CC=C2Cl	
	PCP	*PCP, Angel Dust, Sherman, Sernyl, Wet, Dust, Supergrass, Boat, Tic Tac, Zoom*	Arylcyclohexylamine	1.0–15	-	6	Smoked, oral	C1CCC(CC1) (C2=CC=CC=C2)N3CCCCC3	
Deliriant	Diphenhydramine	*Benadryl*	Ethanolamine	25–700	25–50	6.5	Oral, IM, IV	CN(C)CCOC(C1=CC=CC=C1)C2=CC=CC=C2	
	Mirtazapine	*Avanza, Axit, Mirtaz, Mirtazon, Remeron, Zispin*	Piperazinoazepine	3.5–250	15–60	18	Oral	CN1CCN2C(C1)C3=CC=CC=C3CC4=C2N=CC=C4	
Anxiolytic	Alprazolam	*Xanax*	Benzodiazepine	0.1–3	0.25–2	6.5	Oral, IV	CC1=NN=C2N1C3=C(C=C(C=C3)Cl)C (=NC2) C4=CC=CC=C4	

(Continues)

TABLE 1 (Continued)

Pharmacological class	Drug name	Common and brand names	Chemical class	Non-Medical dosage range (in milligrams)	Therapeutic dosage range (in milligrams)	Average duration of action (in hours)	Route(s) of administration	SMILES	Sources
Antipsychotic	Aripiprazole	*Abilify*	Piperazine	-	2–30	-	Oral, IM, LAI	C1CC(=O)NC2=C1C=CC(=C2)OCCCCN3CCN(CC3)C4=C(C(=CC=C4)Cl)Cl	
	Haloperidol	*Haldol*	Butyrophenone	0.25–10	2–30	24	Oral, IM, LAI	C1CN(CCC1(C2=CC=C(C=C2)Cl)O)CCCC(=O)C3=CC=C(C=C3)F	
	Olanzapine	*Zyprexa*	Thienobenzodiazepine	-	2.5–40	-	Oral, IM, LAI	CC1=CC2=C(S1)NC3=CC=CC=C3N=C2N4CCN(CC4)C	
	Quetiapine	*Seroquel, Xeroquel, Ketipinor*	Dibenzothiazepine	10.0–200	25–800	6	Oral	C1CN(CCN1CCOCCO)C2=NC3=CC=CC=C3SC4=CC=CC=C42	
	Risperidone	*Risperdal*	Benzisoxazole	0.25–6	0.25–6	16	Oral, IM, LAI	CC1=C(C(=O)N2CCCCC2=N1)CCN3CCC(CC3)C4=NOC5=C4C=CC(=C5)F	
Antidepressant	Paroxetine	*Paxil, Seroxat, Loxamine*	Piperidine	-	12.5–75	-	Oral	C1CNCC(C1C2=CC=C(C=C2)F)COC3=CC4=C(C=C3)OCO4	
	Sertraline	*Zoloft*	Tetraline	-	25–200	-	Oral	CNC1CCC(C2=CC=CC=C12)C3=CC(=C(C=C3)Cl)Cl	
	Venlafaxine	*Effexor*	Ethyl-cyclohexanol	-	75–375	-	Oral	CN(C)CC(C1=CC=C(C=C1)OC)C2(CCCCC2)O	

Some testimonials include self-reported age (3,116) and gender (11,129) with which we trained gender-classifying (ROC AUC 0.85) and age-regressing (Pearson correlation 0.56) output heads. Despite missingness, gender class imbalance, and skew toward younger individuals, we can predict age and gender from these reports. Accurate detection of sensitive features is necessary for de-biasing predictions through iterative removal of confounded subspaces (Radhakrishnan et al., 2023; Ravfogel et al., 2020), allowing equitable application of these models in healthcare settings.

Through an entirely different analytic paradigm, BERTowid broadly confirms the salience and ranking of the 11 CCA components (described below). Test set performance predicting the CCA components drops off almost exactly with their ordering by CCA, with (Pearson correlations ranging from 0.68 for CCA 0–0.24 for CCA 11). Expectedly, components explaining more variance are more effectively learned.

3.2 | BERTiment

All emotion predictions generalize to unseen data with ROC AUCs ranging from 0.72 for "Realization" to 0.97 for "Love". Further validation is provided by a hedonic tone classifying BERT model trained with positive and negative movie reviews from IMDB (Maas et al., 2011). The signed Pearson correlations between BERTiment and the valence predictions neatly sorts the fine-grained emotions by hedonic tone. The emotions "Admiration", "Pride", "Approval", and "Love" have the highest positive correlations, while "Annoyance", "Nervousness", "Embarrassment", "Disapproval", and "Disgust" have the largest negative correlations. Originating from entirely different datasets (i.e. movie reviews and Reddit posts) and evaluated on a third orthogonal dataset (i.e. the Erowid testimonials) these models learned mutually reinforcing representations of sentiment, albeit at different levels of granularity, as shown in Figure 1 panel (i).

Domain expert validation for the specific context of the emotions contained in reports of psychoactive experience was provided by a clinical psychiatrist, who manually adjudicated 393 emotions from 256 Erowid excerpts. Concordance between the model and the psychiatrist was within the range of inter-human variability as reported in the original GoEmotions paper (Demszky et al., 2020). Specifically, human labeled emotions were in the top 10 BERTiment emotions for 87% of the labels, in the top 5 for 73%, and in the top 1 for 42%, see Figure 1 panel (k).

Qualitative manual inspection confirms that the extreme (positive and negative) predictions for each sentiment were prominent examples of the emotion (or its opposite), see Supplementary Table 1. As expected with extreme language, profanities, capitalization and modifiers like "very" and "so" are common. Dynamic Time Warping (DTW) quantified the distance between the mean trajectories for each emotion and each pharmacological and biochemical class, in Figure 3 and Supplementary Figure 4 (Berndt & Clifford, 1994). The DTW reveals emotional landscapes that conform to expectations based on pharmacological classifications, molecular

TABLE 2 Drug classifications and testimonial counts.

Drug	Chemical class	Psychoactive class	N testimonials
25i-nbome	Phenethylamine	Psychedelic	150
2C-B	Phenethylamine	Psychedelic	204
2C-C	Phenethylamine	Psychedelic	64
2C-D	Phenethylamine	Psychedelic	47
2C-E	Phenethylamine	Psychedelic	303
2C-I	Phenethylamine	Psychedelic	391
2C-P	Phenethylamine	Psychedelic	58
2c-t-2	Phenethylamine	Psychedelic	118
2c-t-4	Phenethylamine	Psychedelic	16
2c-t-7	Phenethylamine	Psychedelic	173
5-Meo-dmt	Tryptamine	Psychedelic	281
5-Meo-dipt	Tryptamine	Psychedelic	208
5-Meo-mipt	Tryptamine	Psychedelic	71
5-Meo-tmt	Tryptamine	Psychedelic	109
DMT	Tryptamine	Psychedelic	530
DOB	Amphetamine	Psychedelic	45
DOI	Amphetamine	Psychedelic	34
DOM	Amphetamine	Psychedelic	36
DPT	Tryptamine	Psychedelic	148
DXM	Morphinan	Dissociative	457
DIPT	Tryptamine	Psychedelic	54
Ketamine	Arylcyclohexylamine	Dissociative	404
LSD	Lysergamide	Psychedelic	1177
MDA	Amphetamine	Entactogen	79
MDMA	Phenethylamine	Entactogen	1154
Mescaline	Phenethylamine	Psychedelic	136
PCP	Arylcyclohexylamine	Dissociative	76
Psilocin	Tryptamine	Psychedelic	597
Salvia	Salvinorin	Hallucinogen	199
TMA-2	Amphetamine	Psychedelic	24
Alprazolam	Benzodiazepine	Depressant	120
Amphetamine	Phenethylamine	Stimulant	385
Aripiprazole	Piperazine	Antipsychotic	26
Bupropion	Aminoketone	Antidepressant	86
Cocaine	Tropane_alkaloid	Stimulant	568
Diphenhydramine	Ethanolamine	Deliriant	285
Haloperidol	Butyrophenone	Antipsychotic	12
Hydrocodone	Morphinan	Opioid	178
Hydromorphone	Morphinan	Opioid	46
Ibogaine	Tryptamine	Psychedelic	80

(Continues)

TABLE 2 (Continued)

Drug	Chemical class	Psychoactive class	N testimonials
Methadone	Diphenylpropylamine	Opioid	110
Methamphetamine	Amphetamine	Stimulant	419
Mirtazapine	Piperazinoazepine	Deliriant	39
Morphine	Morphinan	Opioid	105
Olanzapine	Thienobenzodiazepine	Antipsychotic	46
Oxycodone	Morphinan	Opioid	237
Paroxetine	Piperidine	Antidepressant	77
Quetiapine	Dibenzothiazepine	Antipsychotic	146
Risperidone	Benzisoxazole	Antipsychotic	31
Sertraline	SSRI	Antidepressant	48
THC	Cannabinoid	Hallucinogen	1348
Venlafaxine	SNRI	Antidepressant	81

structure, questionnaires, and anecdotal reports of drug phenomenology (Preller & Vollenweider, 2018; Studerus et al., 2010). The DTW matrices are skew-symmetric with the sign indicating which drug had higher mean predicted emotion. Figure 4 provides a comprehensive view of the emotional content as determined by BERTiment in the Erowid dataset. Ordered by hedonic-tone, the most negative sentiments are associated with antidepressants and antipsychotics, drugs like opioids and deliriants with both clinical-use and abuse potential are in the middle, and at the positive extreme are psychedelics and entactogens. The association of psychiatric medications with negative emotions is confounded by ascertainment bias of those who seek out these medications, and does not necessarily reflect efficacy.

Zooming in to consider single molecules, MDMA is characterized by both BERTowid and BERTiment in Figure 5. The trajectory of "Love" during MDMA testimonials starts high and ends higher–fitting for a molecule colloquially known as the "love-drug". This arc is clearly distinguished from all other drugs, though closely tracked by the related entactogen, MDA. Supplementary Figure 3 shows the "Sadness" trajectories of stimulants (cocaine, amphetamine, and methamphetamine) are tightly coupled and rise gradually over the course of the testimonial, while antidepressants (paroxetine, venlafaxine, and sertraline) start much higher than the stimulants but gradually fall. In contrast, both stimulants and antidepressants start with similar "Anger" levels, but over the course of the report methamphetamine and cocaine rise while antidepressants increase somewhat less.

The emotions "Realization", "Curiosity", "Confusion", "Surprise", and "Amusement" are consistently elevated in subjective testimonials of hallucinogens and psychedelics as compared to other drugs, most notably the opioids. This constellation of emotions provides discernment within the broad, overlapping classes of hallucinogens and psychedelics as shown in Figure 6. For example, Salvia and DMT are both high in "Realization", and "Curiosity", however Salvia triggers more "Confusion", while DMT generates more "Surprise". PCP is high in "Confusion" and "Amusement", but lower in "Realization", "Surprise" and "Curiosity". The opioids are consistently lower in all of these emotions. "Relief" provides an interesting counterpoint, as it is higher in opioids than in psychedelics, as expected for drugs widely prescribed for their pain-*relieving* effects.

Notably, not every emotion clearly distinguishes different drugs. "Neutral" and "Optimism" trajectories are conserved across pharmacological classes. "Neutral" decreases as the testimonial proceeds, while "Optimism" increases dramatically at the narrative's end, see Supplementary Figure 5. As if the peak-end rule is a self-fulfilling prophecy, testimonials for all drugs end on an optimistic note. The reduction in "Neutral" over the course of a trip is expected as a drug's effects reveal themselves to the user. Notably, a similar reduction in neutral sentiment over the course of the narrative was shown in IMDB movie reviews (Socher et al., 2013).

3.3 | Canonical correlation analysis

We uncovered 11 discrete components that integrate subjective descriptions with the neurotransmitter affinity fingerprint of each drug. These 11 components were statistically-significant after correction for multiple comparisons for all estimated CCA factors ($P < 0.05$, family-wise error-corrected). The explained joint correlation metric, ρ, ranged from 0.86 for CCA 0–0.57 for CCA 10, indicating strong correlation for each factor.

The dominant factor, CCA 0, grouped perceptual phenomena including *visuals*, *colors*, *patterns*, *beautiful*, *light* and *eyes* with apparently transcendental and abstract terms including *reality*, *universe*, *everything*, *ego*, *consciousness*, and *world*. As a gestalt, these terms may be described as "visual-universal". "Visual-universal" terms were linked to the drugs DMT, LSD, and psilocin, primarily serotonergic receptors 5-HT2A, 5-HT1A, 5-HT2C, and localized to the medial prefrontal cortex. The opposite extreme of CCA 0 flagged words suggestive of painful, quotidian ennui, which we summarize as

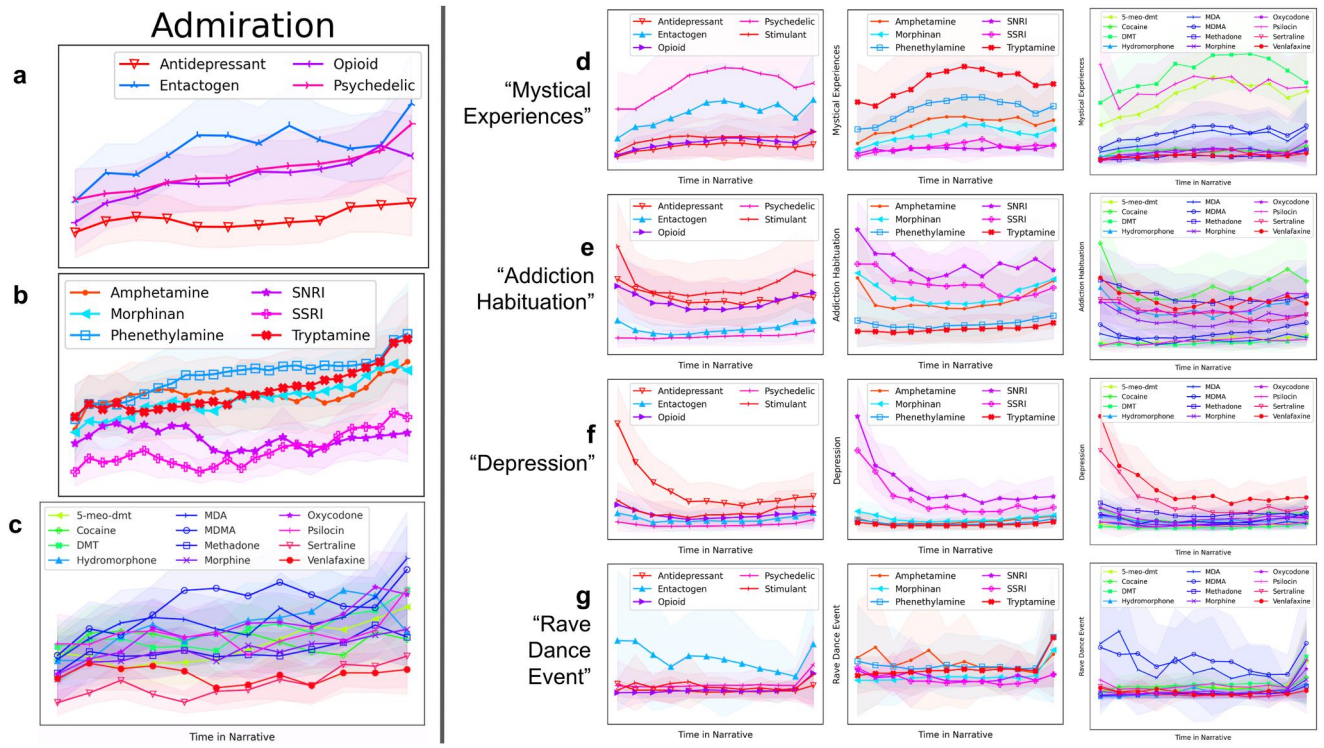

FIGURE 2 Trajectories and Drug Taxonomies. Left: BERTiment trajectories for the emotion "Admiration" at the pharmacological level (a), chemical level (b) and individual drug level (c), see Table 2 for the full drug taxonomy. Note that for clarity we have selected only 12 representative drugs of the 52 included, see other figures for comparisons involving all drugs. Shaded regions indicated +/− intra-drug standard deviation. Right: BERTowid trajectories for each of the 3 different levels of drug classification (from the left panel) on metadata tags "Mystical Experiences" (d), "Addiction Habituation" (e), "Depression" (f), and "Rave Dance Event" (g). Note the concordance between the entactogens MDA and MDMA and the antidepressants sertraline and venlafaxine.

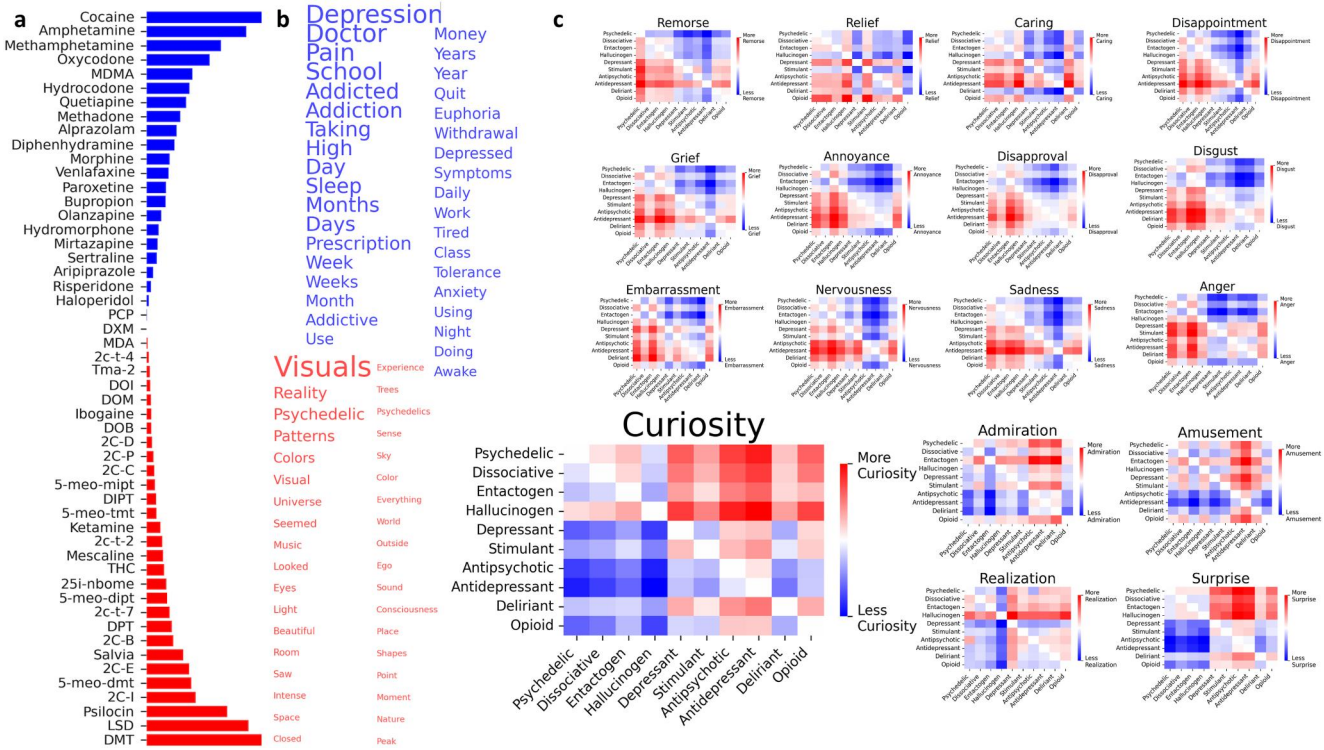

FIGURE 3 Dynamic Time Warp Distances Between 17 Emotions Supports CCA 0. CCA 0, the dominant component from CCA, revealed a ranking of the 52 drugs (a) whose semantic axis, shown by the words at each extreme in (b) were concordant with many emotion DTW matrices from BERTiment (c). Note how all 12 of the DTW matrices in the top half of the figure have bluish upper diagonals and reddish lower diagonals, while the opposite pattern was found for the 5 emotions in the lower half of the image.

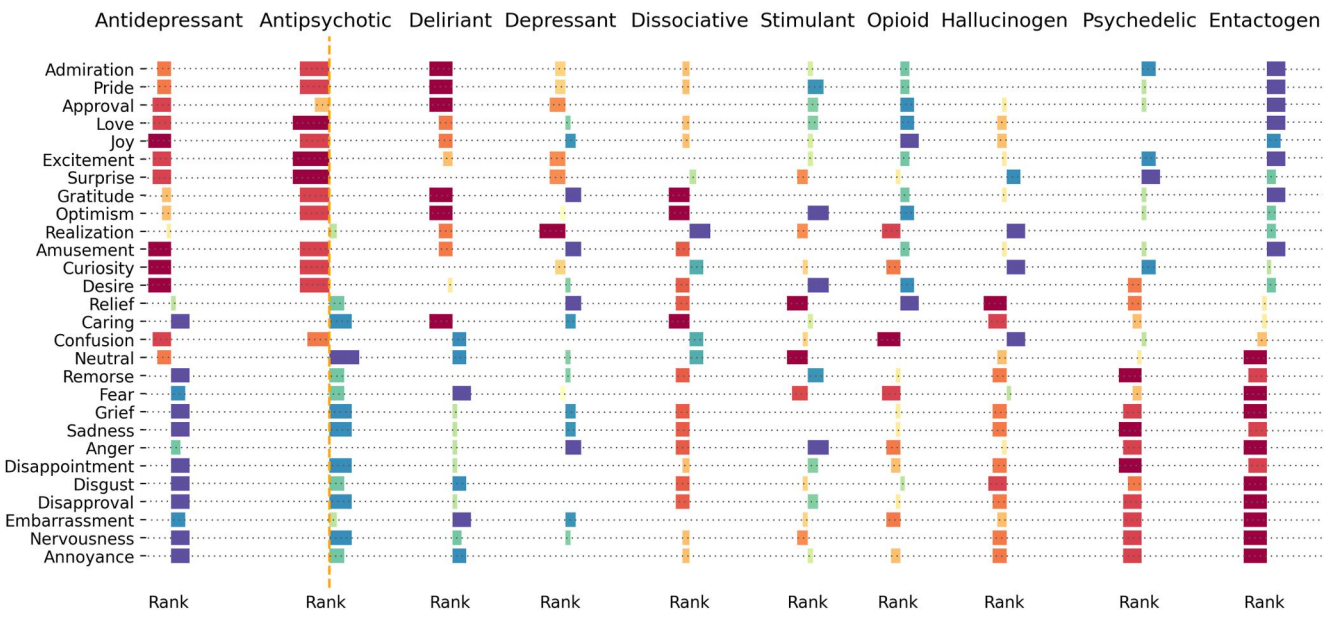

FIGURE 4 Emotionscapes. The relative ranking amongst 10 pharmacological classes for BERTiment predictions for all 28 sentiments. The 28 sentiments are ordered by the hedonic tone spectrum derived from correlation with the IMDB movie review positivity classifier. Note that the drug classes are ordered from left to right according to the prevalence of negative versus positive overall hedonic tone, and that as expected antidepressants show the least hedonic tone and entactogens show the most.

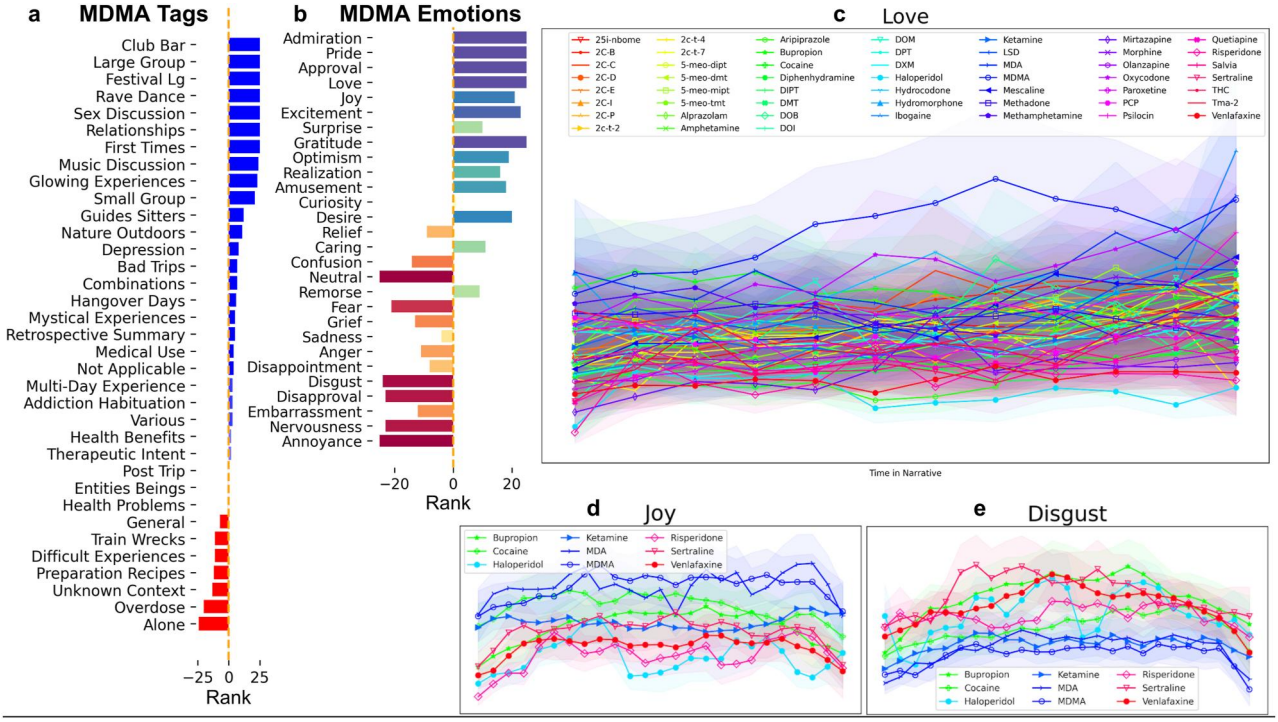

FIGURE 5 MDMA is Singular. (a) The BERTowid tag predictions for MDMA compared to all other drugs, ordered from most associated (blue bars) to least associated (red bars). (b) BERTiment sentiment averages compared all other drugs, ordered by the hedonic tone spectrum, with higher values in blue/green, lower values red/yellow. (c) Trajectory of love compared against all drugs shows that MDMA is the dominant love drug. Sentiment trajectories for "Joy" (d), and "Disgust" (e) for MDMA along with representative drugs from the main pharmacologic classes are shown.

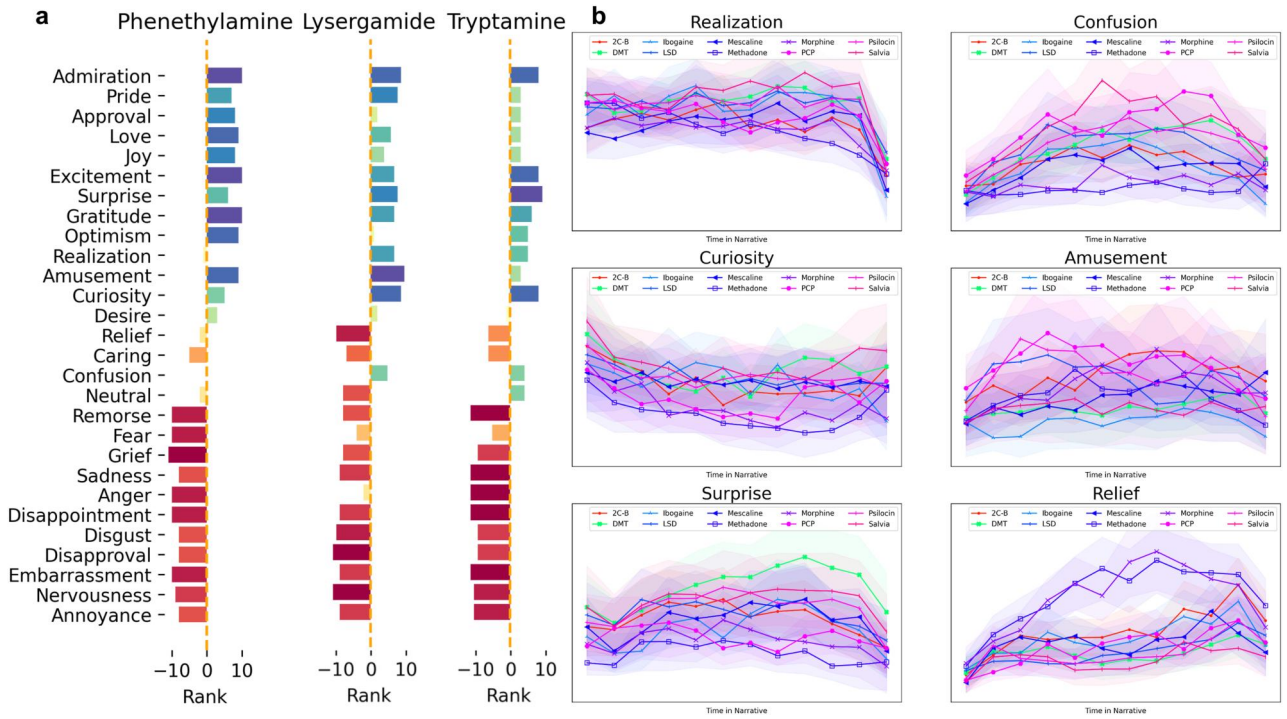

FIGURE 6 Psychedelics are Curious. (a) the three major psychedelic chemical classes are shown with BERTiment's relative ranking of each sentiment, as compared to the 22 classes included in the study. (b) BERTiment sentiment trajectories for a subset of phenethylamines, lysergamide, and tryptamines, and for comparison, opioids and ketamine are also shown. Note how, for tryptamines in particular, the weightings for "Surprise" and "Curiosity" are far greater than they are for MDMA. Whereas "Relief" is much higher for opioids than for any of the psychedelics.

"somatic-suffering": *depression*, *pain*, *addiction*, *sleep*, *awake*, *daily*, *weeks*, *months*, *work*, *and anxiety*. The somatic-suffering pole was linked to cocaine, amphetamine, methamphetamine, and oxycodone, the receptors MOR, DOR, NET, and DAT, and localized to the posterior cingulate cortex and the inferior parietal lobule.

All components are described in detail in Supplementary Figure: CCA 0 - Supplementary Figure: CCA 10. Although not imposed by our analytic model, the cortical mappings of factors often were spatially contiguous with smooth transitions between neighboring regions and bilaterally symmetric between the left and right hemispheres. Additionally, we show that these components are robust to demographic variation in sex and age by showing robust positive correlations between components learned on different subsets of the testimonials, as shown in Supplementary Figure 19.

3.4 | Convergence between BERTiment, BERTowid, and CCA components

Remarkably, the three distinct ML approaches elucidated similar findings. MDMA was found to elicit a unambiguously positive palette of affect, which is of particular interest given its apparent efficacy in PTSD treatment (Mitchell et al., 2021). MDMA was ranked highest among all 52 compounds for "Admiration", "Pride", "Approval", "Love", "Excitement", and "Gratitude". BERTowid tag weightings for entactogens (Supplementary Figures 11 and 12) describe festivity and dancing, but also relational, and ebullient phenomena (e.g. "Glowing"). Similarly, component CCA 5 associated MDMA with *celebration* and *emotional-extremes*.

Results also show important distinctions amongst psychedelic subclasses. In particular, the tryptamines exhibit elevated levels of "Curiosity", "Surprise", and "Realization", while phenethylamines highlight relatively higher levels of "Admiration", "Excitement", and "Gratitude". Tryptamines–especially powerful DMT and 5-MeO-DMT–had higher "Mystical Experiences" tag weightings than phenethylamines mescaline and 2-CE. Interestingly, the chemically distinct diterpenoid compound Salvia shared with these prototypical tryptamines high levels of these emotions and tags though it was also higher on "Confusion" than the tryptamines.

Drug factorization in CCA 0 aggregated stimulants (including MDMA) and psychiatric medications into one pole, associating them with phenomena relating to suffering, addiction, and the mundane. Through an entirely separate analysis, BERTiment DTW demonstrated these drugs all ranked highest for emotions compatible with this theme such as "Disappointment", "Grief", "Annoyance", "Disapproval", "Disgust", "Sadness", and "Anger". The opposite pole grouped psychedelic and hallucinogenic drugs together with terms apparently highlighting a theme of abstract, mystical and beatific phenomena. BERTiment found these drugs to have much higher scores for "Curiosity", "Admiration", "Amusement", "Surprise", and "Realization". This dichotomy between the gloomy, prosaic, quotidian aspects of human sentiment and the more abstract, expansive, and creative

aspects of human potential was the most explanatory distinction in this large corpus of drugs, and the majority of the 28 sentiments considered fit into the same general schema.

The tables of extreme predictions from our models (Supplementary Tables 1 & 2) make clear that our models learn both the positive labels and their opposites. Similarly, Word2Vec demonstrated that directions, not just points, are meaningfully encoded in latent spaces of natural language (e.g. King-Queen = Man-Woman) (Mikolov et al., 2013). Likewise, transformations on our models' representations could encode desired emotional trajectories, for instance separating feelings of confusion from love, or fear from mystical experiences, etc.

4 | DISCUSSION

Shallow CCA and deep transformers, supervised or self-supervised, we have trained models which represent diverse drug experiences in a unified, biochemically-informed, sequentially-sensitive manner. Derived directly from natural language, these representations contain information like the intensity of mystical experience and the depth of joy, anger, and grief. Sequentially applying these models on retrospective reports creates evocative narrative trajectories, which reflect pharmacological distinctions and conform to expectations of subjective effects reported by psychonauts and researchers. Our findings dovetail with prior efforts to use natural language processing tools to analyze the Erowid testimonial database. For example, one recent study (Hase et al., 2022) also found that antidepressants were associated with words denoting negative affect and less with mystical phenomena–which they also partly attribute to ascertainment bias of depressed subjects. A pioneering study from 2012 also noted similarity between pharmacologically-distinct, short-acting drugs like Salvia and DMT (Coyle et al., 2012), while our analysis further distinguished between DMT, linked with higher levels of "Surprise", and Salvia which associated more with "Confusion".

The peak-end rule inspired us to look for trajectories of sentiment, stemming from the counterintuitive finding that there are times when more pain is preferred to less (Kahneman et al., 1993). A rule is not a law, and subtleties like whether pleasure is increasing or decreasing when an experience ends, can be even more important (Mah & Bernstein, 2019). The trajectories produced by BERTowid and BERTiment capture other trajectory shapes which may influence memory consolidation. Combined with clinical outcome data in a prospective manner, this method could potentially help to identify temporally-mediated "rules" correlating clinical efficacy with specific trajectories of experience. There may be a vast repertoire of drug-induced experiential trajectories of clinical import, which are not yet described.

While the trajectories we constructed unfold from the narrative language in the trip report, the methods presented naturally apply to other streams of phenomenological and neurochemical data. Modalities as diverse as EEG, ECG, fMRI, and other biometrics are amenable to trajectory construction, simply by replacing the BERT models with appropriately pretrained encoders (e.g. a 1D CNN for EEGs or ECGs (Khurshid et al., 2022; Liu et al., 2021)). Excitingly, such modalities can be sampled uniformly in time, and independently from the subject's recollection, mitigating uncertainties inherent in self-reported datasets like Erowid about dosage, chronology, and drugs' impact on memory (Doss et al., 2018, 2022). Likely the best representations of acute drug states will be built by multimodally combining subjective and biometric data since both physiological responses and the qualia of experience are critical to understand psychoactive effects.

AUTHOR CONTRIBUTIONS

S.F.F and G.B. conceived of the study and wrote the manuscript. S.F.F. wrote the code to train and evaluate the models, conducted the statistical analyses and created the figures. G.B. annotated the emotions from testimonial excerpts.

ACKNOWLEDGEMENTS

We thank the founders, curators, contributors, and volunteers of Erowid Center for sharing data and for decades of work on the experience report collection.

CONFLICT OF INTEREST STATEMENT

The Authors declare no Competing Financial or Non-Financial Interests.

DATA AVAILABILITY STATEMENT

This study used publicly available data from Erowid (www.erowid.org), the PDSP Ki database (https://pdsp.unc.edu/databases/pdsp.php) and pretrained model weights from TensorFlow Hub (https://tfhub.dev/google/experts/bert/wiki_books/sst2/2). All the code used to prepare the dataset, train the language models and evaluate the conclusions of this work can be found in the GitHub repository: https://github.com/lucidtronix/bertowid.

ETHICS STATEMENT

Human subject research: This work only makes use of previously-released publicly-available anonymized data and does not contain any additional information involving human participants obtained by the authors.

ORCID

Sam Freesun Friedman https://orcid.org/0000-0002-0688-2169

REFERENCES

Altmann. (2000). Memory in chains: Modeling primacy and recency effects in memory for order. *AIB Proceedings*. https://escholarship.org/content/qt7w77c14n/qt7w77c14n.pdf

Ballentine, G., Friedman, S. F., & Bzdok, D. (2022). Trips and neurotransmitters: Discovering principled patterns across 6850 hallucinogenic experiences. *Science Advances*, *8*(11), eabl6989. https://doi.org/10.1126/sciadv.abl6989

Barrett, F. S., Johnson, M. W., & Griffiths, R. R. (2015). Validation of the revised Mystical Experience Questionnaire in experimental sessions

with psilocybin. *Journal of Psychopharmacology*, *29*(*11*), 1182–1190. https://doi.org/10.1177/0269881115609019

Berndt, D. J., & Clifford, J. (1994). Using dynamic time warping to find patterns in time series. Retrieved from https://www.aaai.org/Library/Workshops/1994/ws94-03-031.php

Besnard, J., Ruda, G. F., Setola, V., Abecassis, K., Rodriguiz, R. M., Huang, X.-P., Norval, S., Sassano, M. F., Shin, A. I., Webster, L. A., Simeons, F. R. C., Stojanovski, L., Prat, A., Seidah, N. G., Constam, D. B., Bickerton, G. R., Read, K. D., Wetsel, W. C., Gilbert, I. H., & Hopkins, A. L. (2012). Automated design of ligands to polypharmacological profiles. *Nature*, *492*(7428), 215–220. https://doi.org/10.1038/nature11691

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., et al. (2020). Language models are few-shot learners. Advances in neural information processing systems. https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html

Castaneda, C. (1998). *The Teachings of Don juan: A yaqui way of knowledge*. University of California Press. Retrieved from https://play.google.com/store/books/details?id=BGBYQINGSkgC

Coyle, J. R., Presti, D. E., & Baggott, M. J. (2012). Quantitative analysis of narrative reports of psychedelic drugs. In arXiv [q-bio.QM]. arXiv. http://arxiv.org/abs/1206.0312

Dang, N. C., Moreno-García, M. N., & De la Prieta, F. (2020). Sentiment analysis based on deep learning: A comparative study. *Electronics*, *9*(3), 483. https://doi.org/10.3390/electronics9030483

Demszky, D., Movshovitz-Attias, D., Ko, J., Cowen, A., Nemade, G., & Ravi, S. (2020). GoEmotions: A dataset of fine-grained emotions

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-Training of deep bidirectional transformers for language understanding

Doss, M. K., Samaha, J., Barrett, F. S., Griffiths, R. R., de Wit, H., Gallo, D. A., & Koen, J. D. (2022). Unique effects of sedatives, dissociatives, psychedelics, stimulants, and cannabinoids on episodic memory: A review and reanalysis of acute drug effects on recollection, familiarity, and metamemory. *bioRxiv*, 2022.05.20.492842. https://doi.org/10.1101/2022.05.20.492842

Doss, M. K., Weafer, J., Gallo, D. A., & de Wit, H. (2018). MDMA impairs both the encoding and retrieval of emotional recollections. *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology*, *43*(4), 791–800. https://doi.org/10.1038/npp.2017.171

EROWID: Documenting the Complex Relationship Between Humans & Psychoactives. (2000). Erowid. Retrieved from https://play.google.com/store/books/details?id=rjn4vQAACAAJ

Forgas, J. P. (2011). Can negative affect eliminate the power of first impressions? Affective influences on primacy and recency effects in impression formation. *Journal of Experimental Social Psychology*, *47*(2), 425–429. https://doi.org/10.1016/j.jesp.2010.11.005

Friedman, S., Gauthier, L., Farjoun, Y., & Banks, E. (2020). Lean and deep models for more accurate filtering of SNP and INDEL variant calls. *Bioinformatics*, *36*(7), 2060–2067. https://doi.org/10.1093/bioinformatics/btz901

Galloway, M. (2022). Explaining what causes different psychedelic experiences is a trip. This study shows why. Popular Science. Retrieved from https://www.popsci.com/health/psychedelic-drug-neuroscience-research/

Haertzen, C. A. (1966). Development of scales based on patterns of drug effects, using the addiction Research Center Inventory (ARCI). *Psychological Reports*, *18*(1), 163–194. https://doi.org/10.2466/pr0.1966.18.1.163

Hase, A., Erdmann, M., Limbach, V., & Hasler, G. (2022). Analysis of recreational psychedelic substance use experiences classified by substance. *Psychopharmacology*, *239*(2), 643–659. https://doi.org/10.1007/s00213-022-06062-3

Hasler, F., Grimberg, U., Benz, M. A., Huber, T., & Vollenweider, F. X. (2004). Acute psychological and physiological effects of psilocybin in healthy humans: A double-blind, placebo-controlled dose-effect study. *Psychopharmacology*, *172*(2), 145–156. https://doi.org/10.1007/s00213-003-1640-6

Huxley, A. (2014). The Doors of Perception and Heaven and Hell. Harper collins. Retrieved from https://play.google.com/store/books/details?id=tpkpAgAAQBAJ

Kahneman, D., Fredrickson, B. L., Schreiber, C. A., & Redelmeier, D. A. (1993). When more pain is preferred to less: Adding a better end. *Psychological Science*, *4*(6), 401–405. https://doi.org/10.1111/j.1467-9280.1993.tb00589.x

Khurshid, S., Friedman, S., Reeder, C., Di Achille, P., Diamant, N., Singh, P., Harrington, L. X., Wang, X., Al-Alusi, M. A., Sarma, G., Foulkes, A. S., Ellinor, P. T., Anderson, C. D., Ho, J. E., Philippakis, A. A., Batra, P., & Lubitz, S. A. (2022). ECG-based deep learning and clinical risk factors to predict atrial fibrillation. *Circulation*, *145*(2), 122–133. https://doi.org/10.1161/CIRCULATIONAHA.121.057480

Kubany, E. S., Leisen, M. B., Kaplan, A. S., & Kelly, M. P. (2000). Validation of a brief measure of posttraumatic stress disorder: The distressing event questionnaire (DEQ). *Psychological Assessment*, *12*(2), 197–209. https://doi.org/10.1037//1040-3590.12.2.197

Lawrence, C.-H., Carhart-Harris, R., & Timmermann, C. (2022). Phenomenology and content of the inhaled N, N-dimethyltryptamine (N, N-DMT) experience. *Scientific Reports*, *12*(1), 8562. https://doi.org/10.1038/s41598-022-11999-8

Liu, C., Feng, L., Liu, G., Wang, H., & Liu, S. (2021). Bottom-up broadcast neural network for music genre classification. *Multimedia Tools and Applications*, *80*(5), 7313–7331. https://doi.org/10.1007/s11042-020-09643-6

Loshchilov, I., & Hutter, F. (2017). Decoupled weight decay regularization. In arXiv [cs.LG]. arXiv http://arxiv.org/abs/1711.05101

Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., & Potts, C. (2011). Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the association for computational linguistics: Human language technologies* (pp. 142–150). Retrieved from https://aclanthology.org/P11-1015

Mah, E. Y., & Bernstein, D. M. (2019). No peak-end rule for simple positive experiences observed in children and adults. *Journal of Applied Research in Memory and Cognition*, *8*(3), 337–346. https://doi.org/10.1016/j.jarmac.2019.05.002

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*. *26*. https://proceedings.neurips.cc/paper/2013/hash/9aa42b31882ec039965f3c4923ce901b-Abstract.html

Mitchell, J. M., Bogenschutz, M., Lilienstein, A., Harrison, C., Kleiman, S., Parker-Guilbert, K., Ot'alora G, M., Garas, W., Paleos, C., Gorman, I., Nicholas, C., Mithoefer, M., Carlin, S., Poulter, B., Mithoefer, A., Quevedo, S., Wells, G., Klaire, S. S., van der Kolk, B., & Doblin, R. (2021). MDMA-Assisted therapy for severe PTSD: A randomized, double-blind, placebo-controlled phase 3 study. *Nature Medicine*, *27*(6), 1025–1033. https://doi.org/10.1038/s41591-021-01336-3

Moritz, S., Meier, B., Kloss, M., Jacobsen, D., Wein, C., Fricke, S., & Hand, I. (2002). Dimensional structure of the yale–Brown obsessive-compulsive scale (Y-BOCS). *Psychiatry Research*, *109*(2), 193–199. https://doi.org/10.1016/S0165-1781(02)00012-4

Nichols, D. E. (2016). Psychedelics. Pharmacological reviews. *Pharmacological Reviews*, *68*(2), 264–355. https://doi.org/10.1124/pr.115.011478

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.

Preller, K. H., & Vollenweider, F. X. (2018). Phenomenology, structure, and dynamic of psychedelic states. *Current Topics in Behavioral Neurosciences*, *36*, 221–256. https://doi.org/10.1007/7854_2016_459

Radhakrishnan, A., Friedman, S. F., Khurshid, S., Ng, K., Batra, P., Lubitz, S. A., Philippakis, A. A., & Uhler, C. (2023). Cross-modal autoencoder framework learns holistic representations of cardiovascular state. *Nature Communications*, *14*(1), 2436. https://doi.org/10.1038/s41467-023-38125-0

Ravfogel, S., Elazar, Y., Gonen, H., Twiton, M., & Goldberg, Y. (2020). Null it out: Guarding protected attributes by iterative nullspace projection

Rickli, A., Luethi, D., Reinisch, J., Buchy, D., Hoener, M. C., & Liechti, M. E. (2015). Receptor interaction profiles of novel N-2-methoxybenzyl (NBOMe) derivatives of 2,5-dimethoxy-substituted phenethylamines (2C drugs). *Neuropharmacology*, *99*, 546–553. https://doi.org/10.1016/j.neuropharm.2015.08.034

Roth, B. L., Lopez, E., Patel, S., & Kroeze, W. K. (2000). The multiplicity of serotonin receptors: Uselessly diverse molecules or an embarrassment of riches? The neuroscientist: A review journal bringing neurobiology, neurology and psychiatry. *The Neuroscientist*, *6*(4), 252–262. https://doi.org/10.1177/107385840000600408

Russell, J. A., & Barrett, L. F. (1999). Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology*, *76*(5), 805–819. https://doi.org/10.1037//0022-3514.76.5.805

Salvador, S., & Chan, P. (2007). Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis*, *11*, 561–580. https://doi.org/10.3233/ida-2007-11508 5.

Sanz, C., Zamberlan, F., Erowid, E., Erowid, F., & Tagliazucchi, E. (2018). The experience elicited by hallucinogens presents the highest similarity to dreaming within a large database of psychoactive substance reports. *Frontiers in Neuroscience*, *12*(7). https://doi.org/10.3389/fnins.2018.00007

Sarma, G. P., Reinertsen, E., Aguirre, A., Anderson, C., Batra, P., Choi, S. H., Di Achille, P., Diamant, N., Ellinor, P., Emdin, C., Fahed, A. C., Friedman, S., Harrington, L., Ho, J. E., Khera, A. V., Khurshid, S., Klarqvist, M., Lubitz, S., Philippakis, A., & Westover, B. (2020). Physiology as a lingua franca for clinical machine learning. Patterns. *Patterns*, *1*(2), 100017. https://doi.org/10.1016/j.patter.2020.100017

Schifano, F., Deluca, P., Baldacchino, A., Peltoniemi, T., Scherbaum, N., Torrens, M., Farrě, M., Flores, I., Rossi, M., Eastwood, D., Guionnet, C., Rawaf, S., Agosti, L., Di Furia, L., Brigada, R., Majava, A., Siemann, H., Leoni, M., Tomasin, A., & Ghodse, A. H. (2006). Drugs on the web; the Psychonaut 2002 EU project. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, *30*(4), 640–646. https://doi.org/10.1016/j.pnpbp.2005.11.035

Shulgin, A. T., & Shulgin, A. (1991). *PIHKAL: A chemical love story* (Vol. 963009605). Transform Press.

Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., & Potts, C. (2013). Recursive deep models for semantic compositionality over a sentiment Treebank. In *Proceedings of the 2013 conference on empirical methods in Natural Language Processing* (pp. 1631–1642). Retrieved from https://aclanthology.org/D13-1170

Srivastava, N., Hinton, G., & Krizhevsky, A. (2014). Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine*. https://www.jmlr.org/papers/volume15/srivastava14a/srivastava14a.pdf?utm_content=buffer79b43&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer

Studerus, E., Gamma, A., & Vollenweider, F. X. (2010). Psychometric evaluation of the altered states of consciousness rating scale (OAV). *PLoS One*, *5*(8), e12412. https://doi.org/10.1371/journal.pone.0012412

Sunkin, S. M., Ng, L., Lau, C., Dolbeare, T., Gilbert, T. L., Thompson, C. L., Hawrylycz, M., & Dang, C. (2013). Allen brain atlas: An integrated spatio-temporal portal for exploring the central nervous system. *Nucleic Acids Research*, *41*(D1), D996–D1008. https://doi.org/10.1093/nar/gks1042

Taylor, W. L. (1953). Cloze procedure": A new tool for measuring readability. *Journalism Quarterly*, *30*(4), 415–433. https://doi.org/10.1177/107769905303000401

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in neural information processing systems, *30*. https://proceedings.neurips.cc/paper/7181-attention-is-all-you-need

Williams, J. B. W., Kobak, K. A., Bech, P., Engelhardt, N., Evans, K., Lipsitz, J., Olin, J., Pearson, J., & Kalali, A. (2008). The GRID-HAMD: Standardization of the Hamilton depression rating scale. *International Clinical Psychopharmacology*, *23*(3), 120–129. https://doi.org/10.1097/YIC.0b013e3282f948f5

Yaden, D. B., & Griffiths, R. R. (2021). The subjective effects of psychedelics are necessary for their enduring therapeutic effects. *ACS Pharmacology & Translational Science*, *4*(2), 568–572. https://doi.org/10.1021/acsptsci.0c00194

Zamberlan, F., Sanz, C., Martínez Vivot, R., Pallavicini, C., Erowid, F., Erowid, E., & Tagliazucchi, E. (2018). The varieties of the psychedelic experience: A preliminary study of the association between the reported subjective effects and the binding affinity profiles of substituted phenethylamines and tryptamines. *Frontiers in Integrative Neuroscience*, *12*, 54. https://doi.org/10.3389/fnint.2018.00054

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Friedman, S. F., & Ballentine, G. (2024). Trajectories of sentiment in 11,816 psychoactive narratives. *Human Psychopharmacology: Clinical and Experimental*, *39*(1), e2889. https://doi.org/10.1002/hup.2889