

UC Irvine

2016 Conference Proceedings

Title

Enabling Embodied Analogies in Intelligent Systems

Permalink

<https://escholarship.org/uc/item/0zw4d5hx>

Author

Paolizzo, Fabio

Publication Date

2018-01-08

Peer reviewed



Enabling Embodied Analogies in Intelligent Systems

Proceedings of A Body of Knowledge - Embodied Cognition and the Arts conference
CTSA UCI 8-10 Dec 2016

Fabio Paolizzo

Enabling Embodied Analogies in Intelligent Music Systems

Fabio Paolizzo, PhD

University of California Irvine, Department of Cognitive Sciences and Department of Dance.
University of Rome Tor Vergata, Department of Electronic Engineering.

When human beings make or listen to music, they regularly draw on analogies to other forms of knowledge and embody how they feel through their own body. Research on interactive music systems has gradually shifted from “dialogue” between human and computer to information retrieval, thus posing the risk of commodifying and reifying the encounter with technology [1]. On the contrary, recent advances brought a dramatic increase in human-centered and reflexive retrieval and indexing methods based on subjective concepts such as emotion, preference and aesthetics [2]. Interestingly, recent studies on music cognition suggest that listeners are able to change their semantic relations with the sonic world through functional adaptation at the level of sensing, acting and coordinating between action and perception, in biological, psychological, and cultural terms that involve motor, kinesthetic, haptic and visual, besides the purely auditory components [3]. Implementing some understanding of how we feel and attribute meaning, when interacting with music technology, requires next-generation systems of interactive music, music intelligence and music information retrieval to operate in terms that can represent human cognition as multimodal and embodied.

In the literature of intelligent (and interactive) music systems, early approaches to mutual listening [1] [4] —a milestone in the field—improve with stylistic reinjection; in OMax, a symbolic model recombines the player’s discourse, ‘who is constantly confronted to a reinterpreted version of his own playing’ [5]. Response is similarly contextual in PyOracle, which builds predictive oracles adaptively on arbitrary feature data without quantization, for machine improvisation to match a live musical input in a meaningful way for the user [6]. A step further is taken in interactive reflexive musical systems, which extend the research on active lexicons created from a knowledge model of the user’s interaction [4], for the exploration and manipulation of style [7]. Other systems implement reflexivity by modelling extra-musical information to simulate intentionality in automatic music generation [8]. Still, the gap between computational creativity and human music making is evident, as contrarily to those systems, humans regularly draw on analogies to other forms of knowledge.

Regardless of the widespread interest in computational methods that leverage on the relevance of affect to

1 Lewis, George E. “Too many notes: Computers, complexity and culture in voyager.” *Leonardo Music Journal* 10 (2000): 33-39.

2 Ahn, Hyung-il, and Rosalind W. Picard. “Measuring affective-cognitive experience and predicting market success.” *IEEE Transactions on Affective Computing* 5.2 (2014): 173-186.

3 Reybrouck, Mark, and Tuomas Eerola. “Music and Its Inductive Power: A Psychobiological and Evolutionary Approach to Musical Emotions.” *Frontiers in Psychology* 8 (2017).

4 Pachet, François. “The future of content is in ourselves.” *Computers in Entertainment (CIE)* 6.3 (2008): 31.

5 Lévy, Benjamin, Georges Bloch, and Gérard Assayag. “OMaxist dialectics.” *New Interfaces for Musical Expression*. 2012.

6 Surges, Greg, and Shlomo Dubnov. “Feature selection and live query for audio oracle based machine improvisation.” *Proceedings of the 14th International Society for Music Information Retrieval Conference*. 2013.

7 Papadopoulos, Alexandre, Pierre Roy, and François Pachet. “Avoiding plagiarism in Markov sequence generation.” *Twenty-Eighth AAAI Conference on Artificial Intelligence*. 2014.

8 Paolizzo, Fabio. “VIVO (Video Interactive VST Orchestra): An Interactive and Adaptive Musical Instrument for Self-reflection in Music.” *The international journal of the arts in society* 4.6 (2010): 149-159.

cognition, current research has yet to give greater consideration to the relation between analogy and emotion. We can easily recall three kinds of emotional analogies, such as ‘analogies and metaphors about emotions [...], analogies that involve the transfer of emotions [...], analogies that generate emotions’ [9]. Furthermore, the study of analogical inference regarding emotions should at least take into some account notions of cognition as embodied [10] and the relation between language, visual sequencing and motor planning as information that is similarly structured sequentially [11]. Advancing research for automatic music generation requires enabling systems to classify, recognize and predict emotional states in experiences of ‘musicking’ [12] that comprise domains such as the sonic, visual, motion and linguistic. Investigating how we translate between these domains, “making analogies” is the aim of the Musical-Moods project: a mood-indexed database of scores, lyrics, musical excerpts, vector-based 3D animations, and dance video recordings [13].

Music emotion recognition is an emerging and cross-disciplinary field spanning information retrieval (audio, symbolic, metadata) and machine learning, on a strong backing of music cognition (semiology of music and psychology) and music theory. Strong performance was initially reported in tasks of music mood classification by using audio features only [14], and a number of algorithms are currently available (i.e., MIRtoolbox [15]). In the last few years, other methods incorporating metadata have demonstrated the need for a multimodal approach with a strong focus on human language processing [16] [17]. Recent research has confirmed this by ranking linguistic features higher than audio features, in terms of valence, and by suggesting that symbolic features are useful only when language data cannot be used [18]. Mindprint-based classification is a state-of-the-art technique for language data, which captures sophisticated linguistic features involving semantic, syntactic, and valence information, and notably is not limited to moods leveraged by domain-specific knowledge of interest [19] (i.e., Russell’s [20] [21] or MIREX [8]). Each set of these features, namely a mindprint, is not targeted, as affect features for

9 Thagard, Paul, and Cameron Shelley. “Emotional analogies and analogical inference.” *The analogical mind: Perspectives from cognitive science* (2001): 335-362.

10 Cowart, Monica. “Embodied cognition.” *The Internet Encyclopedia of Philosophy* (2004).

11 Pearce, Marcus T., and Geraint A. Wiggins. “Auditory expectation: The information dynamics of music perception and cognition.” *Topics in cognitive science* 4.4 (2012): 625-652.

12 Small, Christopher. *Musicking: The meanings of performing and listening*. Wesleyan University Press, 2011.

13 Paolizzo, Fabio. *Musical-Moods* (2017): <http://www.musicalmoods2020.org/>

14 Downie, X. H. J. S., Cyril Laurier, and M. B. A. F. Ehmann. “The 2007 MIREX audio mood classification task: Lessons learned.” *Proc. 9th Int. Conf. Music Inf. Retrieval*. 2008.

15 Lartillot, Olivier, Petri Toiviainen, and Tuomas Eerola. “A matlab toolbox for music information retrieval.” *Data analysis, machine learning and applications*. Springer Berlin Heidelberg, 2008. 261-268.

16 Hu, Xiao, J. Stephen Downie, and Andreas F. Ehmann. “Lyric text mining in music mood classification.” *American music* 183.5,049 (2009): 2-209.

17 Turnbull, Douglas R., et al. “Combining audio content and social context for semantic music discovery.” *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2009.

18 Lin, Yi, Xiaou Chen, and Deshun Yang. “Exploration of Music Emotion Recognition Based on MIDI.” *ISMIR*. 2013.

19 Pearl, Lisa S., and Igii Enverga. “Can you read my mindprint?: Automatically identifying mental states from language text using deeper linguistic features.” *Interaction Studies* 15.3 (2014): 359-387.

20 Russell, J. A. “A circumspect model of affect, 1980.” *J Psychol Soc Psychol* 39.6 (1980): 1161.

21 Yik, Michelle, James A. Russell, and James H. Steiger. “A 12-point circumplex structure of core affect.” *Emotion* 11.4 (2011): 705.

emotions would be, and rather spans intentions, attitudes, and emotions. In the study, mindprints are used to represent commonly investigated moods (i.e., anger, annoyance, frustration, boredom, tiredness, calm, satisfaction, pleasantness, joyfulness, etc.), as well as other moods that those broader categories do not capture accurately but, nevertheless, can be expressed through music (e.g., ecstasy, giddiness, craziness, mellowness, etc.).

Research in the Musical-Moods project encompasses a cross-modal approach for validating the model and evaluating the results, through creative practice, in order to realize an innovative music mood classification using audio and metadata. In the last few years, collection of ground truth for mood labels has been increasingly carried out by using “games-with-a-purpose” [22] [23], as based on the “wisdom of the crowds” phenomenon [24] [25] [26] and thus less expensive than listener surveys [27] [28] or social tags [29]. Recently, the approach has been integrated in methods for multimodal distributed semantics, thus showing that ‘the performance of computational models of meaning improves once meaning is grounded in perception’ [30]. The present study proposes a methodology to validate the model through a taxonomy of relations between the musical, linguistic and motion domains. This taxonomy will provide some understanding of how humans “make analogies” in accordance to how they feel, for further exploitation from next-generation human-centered and reflexive, retrieval and indexing technology.

Research methodology and approach

Digital scores inclusive of lyrics will form an effective and large corpus for classification. Such scores are easily gathered from public domain collections of music, converted into audio files by using a software synthesizer, and integrated in a database. Datasets are being collected using a game-with-a-purpose, such as Word Sleuth game-with-a-purpose [31], for Internet users to validate the linguistic model. Furthermore, motion capture technology, such as Vicon systems [32], are being used to track the dancers’ counterpart of physical expression, in order to gather data and further annotations for realizing the cross-

22 Aljanaki, Anna, Frans Wiering, and Remco Veltkamp. “Collecting annotations for induced musical emotion via online game with a purpose Emotify.” *Technical Report Series* 2014.UU-CS-2014-015 (2014).

23 Šimko, Jakub, and Mária Bielíková. “State-of-the-Art: Semantics Acquisition Games.” *Semantic Acquisition Games*. Springer International Publishing, 2014. 35-50.

24 Steyvers, Mark, et al. “The wisdom of crowds in the recollection of order information.” *Advances in neural information processing systems*. 2009.

25 Turner, Brandon, and Mark Steyvers. “A wisdom of the crowd approach to forecasting.” *2nd NIPS workshop on Computational Social Science and the Wisdom of Crowds*. 2011.

26 Lee, Michael D., et al. “Inferring expertise in knowledge and prediction ranking tasks.” *Topics in cognitive science* 4.1 (2012): 151-163.

27 Turnbull, Douglas, et al. “Semantic annotation and retrieval of music and sound effects.” *IEEE Transactions on Audio, Speech, and Language Processing* 16.2 (2008): 467-476.

28 Tsoumakas, Grigorios, Ioannis Katakis, and Ioannis Vlahavas. “Mining multi-label data.” *Data mining and knowledge discovery handbook*. Springer US, 2009. 667-685.

29 Miller, F., M. Stiksel, and R. Jones. “Last. fm in numbers.” *Last. fm press material* (2008).

30 Bruni, Elia, Nam-Khanh Tran, and Marco Baroni. “Multimodal Distributional Semantics.” *J. Artif. Intell. Res.(JAIR)* 49.2014 (2014): 1-47.

31 Pearl, Lisa, and Mark Steyvers. “‘C’moñ–You Should Read This’: Automatic Identification of Tone from Language Text.” *International Journal of Computational Linguistics (IJCL)* (2012).

32 VICON: <https://www.vicon.com/>

modal taxonomy [33] [34][35][36]. In a typical session with the motion capture system, a dancer is asked to: (i) improvise on a music file played from the database and to annotate the mood (e.g., among a proposed set), which s/he considers as expressed by the music; (ii) annotate the mood expressed by the lyrics of another music file in the database; (iii) play the online annotation game. The database then integrates datasets from the motion capture system (i.e., velocity, acceleration, and values of position), their vector-based 3D representation, video signals, and different scores, their audio conversion and lyrics files, for each mood in the model. Third-party computer programs (i.e., MAX/MSP [37]) can be used for carrying out the task. Ground truths and metrics of the human performance on the expression and identification of moods, in music, lyrics and dance, can thus be compared for identifying audio features (i.e., notes per second, attack time, roughness, peak sound level, spectral contrast, MFCCs., CQT, etc.) with some account of symbolic features (i.e., tempo, note duration, note density, etc.). The approach will allow for pre-processing of all the data into mood-indexed excerpts.

Further validation and refinement of the database will be achieved by the realization of music works with interactive and intelligent music systems. This way, case studies will support the identification of features that are best highlighted by music performance or composition practice (e.g., nuances of musical interpretation, cross-domain and analogy-based usage of the datasets, extensive combination of datasets excerpts, superimposition of datasets layers, etc.). In outreach events, the audience will be offered the opportunity to vote for the most interesting works. Finally, a metric of the audience and artists' preference, so generated, will be integrated in the database.

Originality and innovative aspects

The area intersecting interactive and intelligent music systems, information retrieval, as well as computational creativity and music emotion research at large, is rapidly growing and expanding its boundaries, and there is scope for major developments in a short time. Notably, a mood-indexed, multimodal and open-data platform, spanning audio, symbolic, linguistic, numerical, vector-based 3D animations, and dance video recordings, such as that the Musical-Moods project addresses, does not yet exist. In the Musical-Moods project, the database will be ready for further exploitation, thus opening-up new creative possibilities for different types of users, such as musicians, dancers, choreographers, 3D graphics animators, actors, video makers, and researchers operating within these and related areas (i.e., image and video feature extraction and classification). The methodology discussed above is specifically aimed at systems working with various media in these areas. Furthermore, as the way in which humans process and attribute meaning relies on a cross-modal experience of the world, investigating those processes represents a major challenge to the progressive encroaching of technology in content creation, production, fruition and consumption. Thus, a more general objective is to evaluate the contribution that the adoption of elements from the present cross-modal methodology brings to the broader fields of computational creativity, music making and music emotion research.

33 Li, William, and Philippe Pasquier. "Automatic Affect Classification of Human Motion Capture Sequences in the Valence-Arousal Model." *Proceedings of the 3rd International Symposium on Movement and Computing*. ACM, 2016.

34 Dobrian, Christopher, and Frédéric Bevilacqua. "Gestural control of music: using the vicon 8 motion capture system." *Proceedings of the 2003 conference on New interfaces for musical expression*. National University of Singapore, 2003.

35 Shiratori, Takaaki, Atsushi Nakazawa, and Katsushi Ikeuchi. "Dancing-to-Music Character Animation." *Computer Graphics Forum*. Vol. 25. No. 3. Blackwell Publishing, Inc, 2006.

36 Burger, Birgitta. "Move the way you feel: Effects of musical features, perceived emotions, and personality on music-induced movement." *Jyväskylä studies in humanities*; 1459-4331; 215. (2013).

37 MAX: <https://cycling74.com/>

Funding Acknowledgement

The research is supported by the EU through the MUSICAL-MOODS project funded by the Marie Skłodowska-Curie Actions Individual Fellowships Global Fellowships (MSCA-IF-GF) of the Horizon 2020 Programme H2020/2014-2020, REA grant agreement n° 659434.