Semeiosis
SEMIÓTICA E TRANSDISCIPLINARIDADE EM REVISTA.
TRANSDISCIPLINARY JOURNAL OF SEMIOTICS

# Music Recommender Systems: A (Data) Science of Music Aesthetics?

Vinicius de Aguiar
Postdoctoral researcher at the Centre for the Philosophy of Sciences of the University of Lisbon (CFCUL)
vjdaguiar@fc.ul.pt

## Resumo

This paper investigates to which extent and in which ways data sciences and Artificial Intelligence (AI), more specifically recommender systems, are transforming music aesthetics. The interplays between music and AI are not new in history. But while historically, from Athanasius Kircher to Magenta, the focus has been the automatization of creativity, I argue that a new horizon of interplays between music aesthetics and AI has emerged with the massive popularization of streaming platforms powered by recommender systems. I show how the task of music recommendation powered by AI is transforming music aesthetics in three dimensions: epistemological, normative, and phenomenological semiotic.

**PALAVRAS-CHAVE:**
music streaming, music aesthetics, listening interpretant, semeiosis, AI aesthetics

## Abstract

Este artigo analisa até que ponto e de que formas as ciências dos dados e a Inteligência Artificial (IA), mais especificamente os sistemas de recomendação musical, têm transformado a estética musical. As interações entre música e IA não são novas na história. Porém, enquanto historicamente, de Athanasius Kircher ao Magenta, o foco tem sido a automatização da criatividade, eu defendo que um novo horizonte de interações entre estética musical e IA surge com a popularização das plataformas de streaming e seus sistemas de recomendação. Eu analiso como a tarefa de recomendação musical com o auxílio da IA está transformando a estética musical em três dimensões: espistemológica, normativa, e semiótica fenomenológica.

**KEYWORDS:**
streaming de música, estética musical, escuta interpretante, semeiosis, IA estética

Dezembro/2023, v.11 n.1
**ISSN** 2178-5368
**DOI** 10.53987/2178-5368-2023-12-02

## 1 introduction: music and artificial intelligence

The interplays between music and Artificial Intelligence (AI) are not new. In fact, they share a history together. Music's mathematical structures inspired many inventors in the past to conceive combinatorial devices that should be capable of automatically creating music. It is well known the 1842 quote in which Ada Lovelace suggests that the Analytical Engine "might compose elaborate and scientific pieces of music"[1]. Earlier than that, however, polymath Athanasius Kircher, in his *Musurgia Universalis* (1650), had already introduced the *Arca Musarithmica*, a device for composing melodies following purely mechanical procedures.

Since the 50s of the past century, the idea of an AI-composer has become more and more popular and sophisticated, having reached successful commercial applications in recent years[2]. While the quest for originality or "real" creativity animated this research field in its early stages, nowadays, many commercial applications are less interested in proving a metaphysical point or challenging anthropocentric concepts of creativity than in applying sate of the art techniques to generate copyright-free and bespoken music for videos, publicity and even for inducing specific mental states.

In the last decade, however, another type of interplay between music and AI has gained significant prominence. I am referring to Music Recommender Systems (MRSs). As Manovich (2018a, 2018b) has shown, the influence of AI on culture in early 21st century is characterized not so much by the automation of cognition or creativity (e.g. AI-composer), but by the application of algorithms of recommendation on a large, industrial, and global scale, thus "influencing the imaginations of billions" (MANOVICH, 2018a: 2). Manovich (2018b) calls our attention to the role played by "automatic computational *analysis* of the content of all media available online as well as online personal and group behaviors and communication" in the process of shaping contemporary culture (MANOVICH, 2018b: 474). Arguably, AI has disrupted musical cultures much more radically since it entered the domain of recommendations than since it started to generate music on their own. But while data scientists have been actively discussing the challenges and opportunities of MRSs, their consequences for music aesthetics are still poorly understood. In the next sections, I will argue that *MRSs are characterized by complex interplays between sciences, technologies, and aesthetics that demand further clarifications*. This paper aims to contribute to this clarification.

## 2 on the scope of music aesthetics

How are MRSs are transforming music aesthetics? First, some notes on the scope of music aesthetics. One can define music aesthetics as a domain of philosophical aesthetics or as a domain the philosophy of music, having its roots in the works of authors such as Baumgarten and Hanslick[3]. This philosophical approach already implies a wide range of questions beyond the topic of beauty in music, such as perception, values, meaning, ontology, and metaphysics of music. However, we can expand the scope of music aesthetics by adopting the definition of Fubini, according to whom music aesthetics is "any type of reflection about music, its nature, its ends, and its limits" (FUBINI, 2008: 14, our translation). This would include also theories traditionally known as musicological, such as those related to music analysis, as well as philosophical, sociological and even scientific approaches to music. In this linage, music aesthetics is presented as an effort to understand (aspects of) music. In this vein, music aesthetics seems to have an eminently theoretical vocation.

But besides that theoretical dimension, I believe we can also understand music aesthetics with reference to more practical elements, such as styles, genres, creativity, performance, artifacts, perceptions, judgements, taste, and all sorts of doings of and through music. This is what philosopher and media scholar Dieter Mersch (2015) calls the epistemologies of the aesthetics, that is, the modes of knowing through aesthetic practices such as composing, performing, and listening. Though traditionally music aesthetics is concerned with discourses about music, I believe it is important to take into account also these “musical ways” of engaging with music. It is important to acknowledge that such situated enactments of and with music generate their own type of musical knowledge that also participate in the broad field of music aesthetics, for example in the form of aesthetic choices, aesthetic methods, aesthetic models and so on. In de Aguiar (2022), I have analyzed various cases of such musical epistemologies focusing on their pragmatic and semiotic — to be more specific: on their diagrammatic — dimensions.

Historically, the interplays between the scientific or theoretical aesthetics and the practice-oriented aesthetics have been many. Sometimes, the same individual developed both aesthetic domains in dialogue with each other, such as Jean-Philippe Rameau. But in most cases, what happens is a back-and-forth movement in which aesthetic processes and artifacts inspire theories and vice-versa.

Moreover, it is possible to determine many intermediaries between theoretical and empirical music aesthetics, such as models, technologies, and habits (see Fig. 1). One the one hand, theoretical knowledge can become operational models and be implemented in technologies. On the other hand, there are habits that emerge from individual or socially shared practices, thus becoming the object of theoretical music aesthetics. Evidently, the use of certain technologies in musical practices, precisely because such technologies might be theoretically informed, is one way in which theoretical knowledge pours into musical artifacts or processes.

To make things even more complex, theories, practices, and technologies, in the context of music aesthetics, might also receive inputs from multiple areas, some of which might have little to do with music in a first moment. For example, the resolution of a mathematical problem (the twelfth root of two) was fundamental to inform musical instruments and musical practices, leading to the tonal aesthetics and all theories and practices related to it. Likewise, the birth of modern computers and the Internet transformed musical media and eventually led to a radically new socio-cultural ecology of music with streaming platforms.

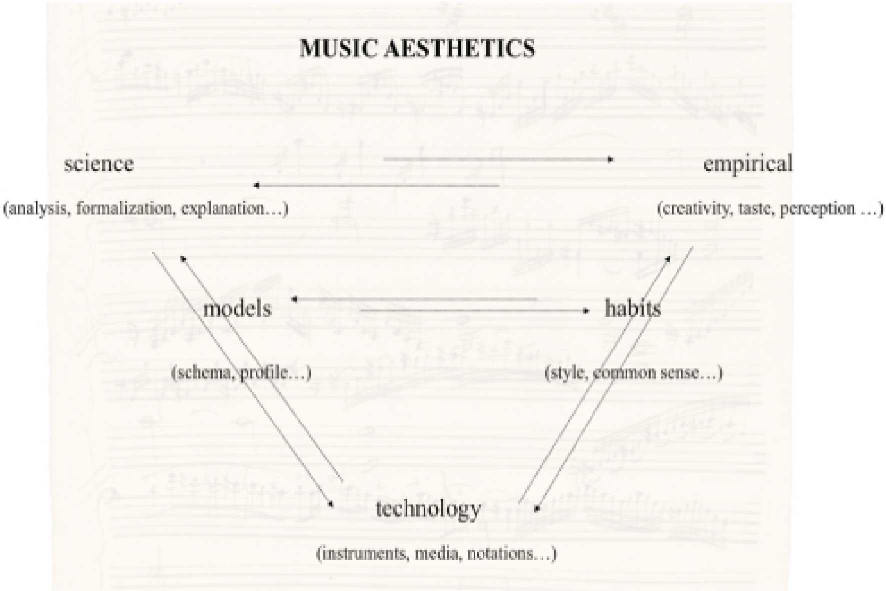


**Figure 1.** Diagram of the expanded field of music aesthetics.

It is probably possible to identify many variations and subdivisions within the above diagram. Be it as it may, for the purposes of this paper, that framework should suffice to guide our analysis of this technology – which is also a science, a model, and a musical practice – that has arguably disrupted and is currently reshaping music aesthetics: music recommender systems.

## 3 AI and Music Aesthetics: The Case of Music Recommender Systems

In the specialized literature, MRSs are defined with reference to the so-called information overload problem. As the story goes, since the rise of digital music distribution, "music listeners are suddenly faced with an unprecedented scale of readily available musical content, which can easily become burdensome" (SCHEDL et al., 2022: 928). It is in this context that recommender systems are developed in multiple domains to provide suggestions and support users in decision-making processes (RICCI, ROKACH, BRACHA, 2022: vii).

Research on MRSs has been attracting considerable attention from data scientists. Multiple types of MRSs keep being proposed to address this challenging task. Schedl et al. (2022) and Schedl et al. (2015) provide extensive and thorough presentations of how MRSs function in computational terms. Based on their paper, we could say that, roughly, MRSs are characterized in terms of the input data available, techniques used to process the data, and tasks.

Common examples of input data are explicit and implicit feedback provided by users. In general, MRSs are fed with data coming from users' ratings to items (e.g. song or artist) as well as binary feedback such as Like or Favorites. Systems can also leverage from user-related information, such as those provided via semantic queries (e.g. activity, emotional state, etc.). Other specific sources of input data for MRSs include the history of listening events, play counts, total time listened (to songs, albums, artists, etc.), and track skips. Besides that, MRSs can also leverage from musical metadata such as manual annotations and social tags. Finally, analyses of audio content, which go under the domain of Music Information Retrieval, are also an important source of input data for MRSs.

Examples of specific tasks are the prediction of user ratings for specific parameters such as artists, album, track (rating prediction task) and the prediction of whether a user will consume a specific item (prediction of item consumption behavior). Tasks can also be distinguished between lean-back recommendations, in which user interaction with the system is minimum, and lean-in recommendations, in which the goal is more exploratory, usually mediated by semantic queries.

Regarding the techniques employed in MRSs to process data and successfully complete the task, they are plenty, as we can see in overviews such as Schedl (2015). Generally speaking, they involve both general data science methods (e.g. process mining, statistics) and more strictly speaking AI methods (e.g. machine learning algorithms)[4] that are employed to process the input data available and generate models. However, it is possible to identify some recurrent types.

Collaborative-Filtering (CF) is a classic approach to the recommendation problem. In CF approaches, input data comes from user-item interactions (implicit or explicit). The underlying idea is to leverage from users with similar profiles to predict whether a user will listen to a certain previously unseen item. If user A has liked item B, and if user X has a similar profile than user A, then, it is likely that user X will also like item B. Another classic approach is known as Content-Based Filtering. In this type of approach, algorithms are designed to match content information for each specific user. This means if user A likes item B, and item C is similar to item B, then user A is likely to enjoy item C as well. Within CF and CBF, many different types of algorithms can be employed. Moreover, there are also plenty of hybrid approaches, not to mention more recent developments in terms of algorithms that also use contextual and user-related information to predict ratings and relevance of new items for each specific user.

Less understood, however, is the relationship between MRSs and music aesthetics broadly construed (see Section 2). In what follows, I will conduct this analysis.

An important clue in this direction can be found in Arielli's "simple map" of how AI is related to aesthetics (MANOVICH & ARIELLI, 2021: 11-15). Arielli introduces the following important distinction: current computational approaches to culture deal with objects and subjects as well as with abstraction and generation. Applying this map to MRSs, I would say that they are technologies that analyze not only songs (*objects*) but also music consumption behavior and experiences (*subjects*). Moreover, these technologies are not interested only in extracting features and patterns (*abstraction*) but also in generating outputs such as playlists and aesthetic experiences (*generation*).

| | **Pattern recognition** (analysis and description) | **Pattern generation** (production and prediction) |
|---|---|---|
| Objects | **Studying objects** | **Generating objects** |
| Subjects | **Studying subjects** | **Generating subjects** |

**Figure 2.** Arielli's "simple map" of AI-Aesthetic relationships. In Manovich & Arielli (2021: 12).

Following those clues, my analysis of MRSs will focus on: (i) MRSs as science; (ii) MRSs as generative model; (iii) and MRSs as experience. To put it in other words, in what follows, I will analyze the epistemic, normative, and phenomenological roles played by MRSs in the domain of music aesthetics (see Fig. 1).

### ***epistemology: music recommender system as science***

MRSs allow us to diagrammatize music in at least three fundamental layers. For the sake of simplification, let us call those layers by (i) audio content (information contained in the audio files), (ii) music metadata (such as release date and album cover), (iii) and user-related content (e.g. feedback, consumer behavior). Following a Peircean jargon, we could say that their objective is to identify habits (i.e. generalities, patterns) and convert them into computer-generated diagrams, such as extensive lists, tables, graphs, models that are supposed to represent the knowledge acquired and afford further manipulation of it (e.g. processing, filtering, grouping, predicting). Therefore, we can say that MRSs are epistemic and cognitive tools.

In this sense, MRSs can be described as a (at least potentially) new scientific approach to music, not so distant from mathematical and computational theories of music. As is well known, the epistemic power of mathematical formalizations to understand how music "works" has been widely exploited since the Pythagoreans (MAOR, 2018). And the emergence of computational approaches to music is in no way alien to this tradition. In a way, we could say that AI builds on this tradition. AI has added to that linage the technological — and, through that, the epistemological — power to make analyses on a different scale (not analyses of individual songs or pieces but of whole datasets composed of millions of songs) and about new objects: the target is no longer audio features only, but also the relation between audio, metadata, and user behavior.

It is important to notice that, so far, it is mainly the demands and goals of the industry that have been shaping the research on MRSs. Certain fundamental notions are not questioned or reviewed regarding their epistemic validity. On the contrary, they are presupposed as valid and necessary parameters in the design of MRSs. In this field, machine learning algorithms are meant to detect in their datasets aesthetic concepts such as "relevance", "preference", "personalization", "similarity", "novelty", and ethical notions such as "fairness" and "bias". They are meant to perform epistemic procedures such as "prediction". And they are supposed to be evaluated regarding their success in performing the task of music recommendation. Needless to say that all those notions and procedures receive a mathematical formulation in order to be dealt with by computers.

However, one could ask: to which extent such a conceptual framework thus defined in mathematical terms is adequate for a scientific understanding of aesthetic relevance or aesthetic experiences? In other words: to which extent is this framework adequate to generate new knowledge about what makes a song relevant or interesting, in general or to a specific user? Isn't this method being employed less as a research method and more as a sort of *administrative technique*, to use Adorno's criticism of the Radio research (ADORNO, 2009: 134)? Are such systems being designed and used not so much to acquire new knowledge about music and taste but rather to identify patterns that can be exploited according to commercial standards? Isn't the current interest of the MRSs research community less a matter of using their tools to develop a new science of music aesthetics and more a question of using data sciences to achieve certain commercial results, such as "satisfying" the user by "adapting recommendations based on intrinsic user characteristics" (SCHEDL et al. 2022: 960-1)?

Criticism to MRSs coming from the humanities have pointed some questionable implications of those technologies, such as the impact of personalization in musical sociality (HANRAHAN, 2018) and the influence of recommendation algorithms and other metrics in the economy of musical distribution (MAASØ & HAGEN, 2020). However, I believe that we need to stress more the *epistemic limitations* of MRSs as currently developed and implemented and how to overcome them. As Prey (2018) has argued, we could say that the way platforms are *conceiving* their listener is less a matter of trying to *understand* them and more a question of *generating* a certain type of listener.

Given those limitations, an alternative approach is to use the same methods of MRSs but *less biased by commercial imperatives*. We do not have to assume a "pure" epistemology in which methods such as those borrowed from data sciences would disclose the "real" listener, their taste, preferences, interests and so on. Instead, and to give just a couple of examples, MRSs as a scientific method *could be used more systematically* to study the historical transformations of digitalized music in terms of acoustic properties (e.g. SERRÀ ET AL., 2012; CHINOY & MA, 2018) or the transformation of consuming patterns of music fans.

In sum, though MRSs can be described and could function as an epistemic tool for a new paradigm of music aesthetics, up to now, the interest in this method has been reduced to the identification of patterns relevant for commercial purposes and in using the *provisory knowledge* acquired as a tool to generate norms and experiences.

### *normativity: music recommender system as generative model*

As was briefly mentioned in the beginning of this paper, the figure of the AI-composer dominated the debate about AI and music for some decades. Less studied, however, is the role played by AI and data science techniques through recommender systems in the generation of new music and new musical artifacts.

The influence of scientific models on musical composition is not new in history. The Pythagorean metaphysics was fundamental to determine the tuning of instruments, musical scales, the hierarchy between tone intervals and so on. The physics of resonance and overtones played a central role in the systematization of tonality. I believe that nowadays, data-driven knowledge about music and users acquired through MRSs is playing a role in shaping how new music is designed. In fact, I believe they are playing a *normative* role in shaping the aesthetics of new music. This can be attested by analyzing a typical musical product of the age of streaming: the playlists.

Techniques employed by MRSs certainly provide insightful metrics for musicians/composers — the research field known as the Hit Song Science illustrates that very well (e.g. DIMOLITSAS, KANTARELIS, FOUKA, 2023). However, playlists are more clearly generated based on knowledge acquired through MRSs. The analyses of audio data, music metadata and user-related data that characterize MRSs provide the relevant metrics to choose which songs should figure in which type of playlists and for which user/context. For example, these technologies are responsible for predicting which songs a certain user is likely to consume but also which songs should figure in a certain thematic, mood-oriented, playlist.

When we consider the case of automatically generated playlists, there is almost a feedback loop between dataset, information processing, model generation, and the personalized playlist. However, even assuming the intervention of a human curator, the process of designing the playlist is not indifferent to the knowledge acquired through MRSs. Bonini & Gandini (2019) call this process as algo-torial, a sort of augmented or enhanced editorial power.

Be it as it may, the fact is that playlists are designed following criteria first identified through techniques such as those operating in MRSs. I believe we can call these technologies a kind of *technological enhancement* of the traditional use of statistics in the music industry. Nowadays, besides traditional parameters such as sales targets and success ratios, technologies such as MRSs are capable of tracking and processing many more intimate and continuous sources of data in order to identify relevant patterns to be further exploited in the design of new, bespoken, musical outputs (cf. OSBORNE & LAING, 2023).

### *phenomenological semiotics: music recommender system as experience*

As we have seen, the new methods provided by data sciences and AI for the analysis of music and music-related data, and more specifically the case of MRSs, impacts not only the epistemological dimension of music aesthetics, but also to its practical or normative aspects such as creativity and aesthetic outputs. But how about the aesthetic experience of music: is it also influenced by the methods of recommender systems? And if so, how?

In the past years, streaming platforms have been criticized by academics and media for promoting functional, bland, unchallenging, passive, distracted, and fragmented music aesthetic experiences (cf. HESMONDHALGH, 2022). Personalization is usually considered a key variable to understand the impact of music streaming on the aesthetic experience of music. Moreover, based on our previous section, it is obvious that these technologies influence the aesthetic experience of music to the extent that they provide models for composers and curators to design new musical products. However, I believe there is more to it and that MRSs have become fundamental not only in the design of listening stimuli (e.g. songs, playlists) but also in the design of listening experiences more directly. I believe that the semiotic character of those experiences can help us better understand the role of MRSs in the constitution of music aesthetic experiences.

The semiotic dimension of music has been widely debated since Peirce (e.g. NATTIEZ, 1990; MONELLE, 2010; TARASTI, 2012). A summary of the main arguments goes much beyond the scope of this article. My goal is simply to *use* parts of the Peircean scholarship to analyze the questions aforementioned.

The first thing I would like to point out is the semiotic character of music aesthetic experience. Following Peirce's semiotic phenomenology as recently reconstructed by Paolucci (2021, pp. 81-9), I would argue that instead of assuming that music aesthetic experience is something immediate or radically subjective, it is important to start by acknowledging that, here too, takes place an operation of and through signs, i.e., a semeiosis [5]. Music aesthetic experience constitutes a phenomenon, it is "something which appears", and thus it is analyzed through the phenomenological method. But, this appearance is also an effect, something that is elicited by "something else"; therefore, it is fundamentally a sign. Finally, as a sign, music aesthetic experience is not static, but a dynamic process temporally distributed[6].

As a type of sign, music aesthetic experiences should not be characterized in terms of internal/external, subjective/objective, but in terms of sign relations and sign actions, i.e., semeiosis. In fact, this semiotic dimension of music aesthetic experience help us understand how fundamentally linked it is to "external" signs, such as material artifacts, intersubjective interactions, social and cultural habits and so on. This semiotic ontology of the aesthetic experience of music also helps to clarify what Flusser (2014) explained saying that "[...] one listens differently, depending on whether one is hearing opera or Indian ragas, and opera differently, depending on whether the opera is on television or on vinyl" (FLUSSER, 2014: 112). This is so precisely because the aesthetic experience of music is a semiotic effect of this "something else" that it represents; and when this "something else" is materially transformed, the effects it fosters in the interpretant mind are also bound to change.

What kind of sign is the aesthetic experience of music? I propose to define the aesthetic experience as an interpretant sign. For the sake of simplicity, I propose to call it the *listening interpretant*. As such, the aesthetic experience of music is the effect of a sign that represents to the listening interpretant a certain object. I would argue that, in this triadic relation, the sign is constituted by the music performed within a multimodal situation, such as a record player in a room or an orchestra in a theater. The object of this sign would be an aesthetic experience of music as arranged by the composer (and eventually by other intermediary parties involved such as a conductor, producer, musician, etc.). The sign, in this case a particular material instantiation of the music aesthetic experience as conceived previously by the composer, fosters a listening interpretant, which would be a case of music aesthetic experience[7]. This listening interpretant will, of course, become the semiotic object of further semeioses. This is so because a listening interpretant can be further analyzed, revised, improved, corrected. This amounts to saying that the listening interpretant is subjected to (logical) self-control.

It is also worth noting that the example of music in Peirce's semiotics has often appeared associated with features such as iconicity and emotional interpretants (cf. SHORT, 2007: 204; BELLUCCI, 2017: 326; see also PEIRCE, CP 5.475). However, it has also been pointed out by many authors that the semiotic status of music is much more complex and, in different degrees, it can cover the whole spectrum of phenomenological and semiotic categories identified by Peirce (e.g. CUMMING, 2010; KRUSE, 2007; OLIVEIRA, 2010; SANTAELLA, 2018; COLAPIETRO, 2020; ATÃ & QUEIROZ, 2021). That is why I have extensively argued elsewhere (DE AGUIAR, 2022) that music listening has the form of a diagrammatic process in which a first schema apprehended, be it too vague or too structured, can be further revised and transformed thus deepening the attunement with and ultimately the understanding of the musical meaning.

How do MRSs participate in the semeiosis of music aesthetic experience as described above? It is interesting to notice that, although MRSs can play a normative role in the design of musical signs, they are not musical signs in the strict sense; they are not themselves part of the arrangement of tones that constitute a song. They are also not exactly the medium through which the musical sign is conveyed.

I would argue that an overlooked role played by recommender systems in the semeiosis of music is that of a *selector of signs*. The primary function of a MRS is to select which musical sign will be presented to a certain mind given certain variables (input data) and aiming at certain effects. By doing so, those technologies play a direct role in establishing which musical signs will be present in one's horizon of consciousness or field of attention at a given moment.

We could explain this with reference to Peirce's notion of sheet of assertions. As Zalamea has summarized, "[t]he realm of Peirce's *continuum* is represented by a blank sheet of assertion where, following precise control rules, some cuts are marked, through which information is introduced, transmitted and eliminated." (ZALAMEA, 2001: 73). Now, this notion of a continuum represented by a blank sheet of assertion was recently generalized by Stjernfelt (2022) to cover not only the domain of the Existential Graphs, but also the whole field of perceptions and thoughts "in the wild". As noticed by Stjernfelt, "[i]n the wild, we find a lot of different cases where "Sheets", that is, *delimited areas of attention within the field of perception*, have the same function: signs placed on such a Sheet are cognized as fused together into one asserted proposition" (STJERNFELT, 2022: 110, my emphasis). Going beyond the examples of logical notation and linguistic propositions, Stjernfelt argues that there are other modalities for joining subjects and predicates or even multiple propositions within a "sheet of assertion" as to convey meanings. Fundamentally, this semiotic operation is characterized by the spatiotemporal co-localization of signs, such as in maps, diagrams, book covers, posters and so on (STJERNFELT, 2022).

Following those insights, I would argue that the output of a recommender system is like a sign drawn on a *temporal sheet of assertions*. It constitutes a kind of affirmation, a positive statement positing that a certain sign (i.e., a song) is the case given the circumstances (e.g., certain user, situation, or goal). Moreover, this sign takes the form of a subject to which predicates are spatially connected within the platform (e.g., title, artist, comments, images) conveying a multimodal proposition[8]. This would be the most fundamental semiotic role of a MRS. Evidently, the semeiosis becomes much more complex if we consider the interplays between multiple sign/song recommendations that are co-localized spatially *and temporally* with their respective predicates co-localized with them on a 2-dimensional surface. In a concrete example such as opening the Spotify main page, we could say that the multiple selected signs with their respective predicates are in a dialogical relationship with each other. Their co-localized presence on this sheet of assertions conveys a specific meaning or entails certain effects that the signs/songs isolated would not foster. An analysis of this complex semeiosis goes much beyond the scope of this paper and should be developed in further research.

This sheet of assertions I am referring to constitutes not so much the physical 2-dimensional space of the platform/screen, but rather the field of attention of the listener. The algorithmic part of the process, i.e., the recommender system, constitutes the control rules or the rules of transformation of the possible assertions; they determine whether a sign (e.g., a song) is “true” or “false” given a certain user and given a certain objective. The input data constitutes a dynamic and growing Universe of Discourse of the recommender system. Finally, we could consider the input data as itself a sign of an even wider Universe of Discourse which would consist in the whole of music and musical-related semeiotic objects (e.g., songs, writings, and experiences) that transcends a certain dataset.

MRSs are responsible for the design of sheets of assertions composed by certain signs and not others. Because of that, music recommender systems play a central role in shaping the listening interpretants, i.e., the aesthetic experience of music. They choose the signs (e.g., individual recommendations), and in some cases even the temporal order/hierarchy of the signs (e.g., playlists), that will be present to a certain listener at a given moment. By doing that, they determine, at least partially, the meanings, that is, the effects to which a listener will be subjected to at a given period of time. In this sense, these algorithms exercise a certain control over aesthetic habits, over habits of feeling music, and even over emotional experiences aided by music.

As a semiotic process, i.e., semeiosis, the presentation of certain signs will not generate effects in a mechanical way. If that were the case, all listeners presented to the same signs would have the exact same experiences. The listening interpretant will also depend on previously acquired habits. After all, the interpretation of the same sign can vary significantly depending on the sensibility of the interpreter to different features conveyed by that sign. Still, albeit the influence of habits in the constitution of listening interpretants, the aesthetic experience is not a matter of pure subjectivity, pure constructionism, or a mere internal phantasy of one’s imagination, so to speak. The role of the sign presentation should not be neglected. Signs embody meanings which means, quite pragmatically, that they affect the conduct of those that engage with them. And while the degree of this affection might be dynamic and thus subjected to revision and transformation (not to say: correction), the emotional experience is not at all arbitrary.

A semiotic critique of MRSs would ask questions such as: Why certain signs were selected for the sheet of assertion instead of others? What is the pertinence of the logic of the selection employed by a certain system? What is the criteria of relevance and what validates it? How does the system calculate the meaning of each sign and how accurate this calculation is? Can the system promote the transformation of listening interpretants, namely the abduction of clearer listening interpretants?

If we consider the commercial applications of MRSs, we could speculate that due to commercial interests, the promotion of meanings whose effects can be calculated with a higher degree of precision tend to predominate, thus favoring the crystallization of certain aesthetic habits. Take the case of mood-related playlists. Even considering the intervention of human editors, it is safe to assume that a great deal of the selection process of signs to influence the mood of the listener is calculated or decided by algorithms. We have to ask, therefore, exactly what kind of effect is being calculated by algorithms whose task is to recommend songs that are "sad", "happy" or whatever other mood, and how exactly is such calculation done. Then, we have to consider the case of moods that are not so self-evident and of songs/signs whose meaning is far from being exhausted, songs that are ambiguous, for example. Could it be that those signs will simply not be promoted? Could it be that they will be promoted within a context that does not do justice to their meanings? As we can see, the use of algorithms to select musical signs according to the intended meaning or effect they are likely to generate in the listener is highly challenging. Therefore, there is the risk of reducing the variety and complexity of listening interpretants through the promotion of few schematic types that are easier to calculate.

It is also important to notice that, given the commercial and other practical interests, the knowledge used to calculate the meanings, i.e., the effects of different signs — in this case, to calculate the expected listening interpretant — is only loosely subjected to the pragmatic rule of self-controlled revision and criticism. Currently, it seems to me that research on MRSs (e.g., GOTO, 2012; SCHEDL ET AL., 2014; SCHEDL ET AL., 2022) is less interested in using algorithms to enrich our *understanding* of the complexity of listening interpretants and more interested in identifying in them measurable patterns that can be efficiently exploited in the context of predictive models, *even if these patterns represent only a provisory and partial knowledge about the aesthetic experience of music.*

Now, the problem with using provisory and partial knowledge to select which musical signs and therefore which musical meanings or listening interpretants to promote is that this might risk having a negative impact in the aesthetic experience of music, e.g., failing to uplift the aesthetic experience of music or failing to generate new types of interesting experiences. This is so because, although every knowledge can be considered provisory (cf. Peirce's fallibilism), in certain cases, the knowledge *chosen* to be applied and the *means* of implementations of such knowledge might simply not be up to the task at hand. To illustrate this risk, let us consider the following cases. We know that there is music with lyrics and music without lyrics. However, if we were to apply this criterion and this criterion only, this knowledge would not help us to make much sense of a large musical dataset. Consider, now, the common-sense knowledge according to which there are happy and sad songs. While it is true that it is possible to use algorithms to identify features that are likely to be associated with the affective dimension of songs, it is also true that such algorithmic mapping at a certain stage of this development might simply not be enough to provide a richer picture of musical affects. Therefore, if we decide to apply such a provisory knowledge to help millions of people navigate through enormously large datasets of music, we might risk making music aesthetic experience a rather uninteresting business.

## notes

Acknowledgements: I would like to thank the two anonymous reviews for their comments on an earlier version of this paper. This work was financed by national funds through FCT - Fundação para a Ciência e a Tecnologia, I.P/MCTES, under the Research UIDP/00678/2020 and UIDB/00678/2020.

[1] A full transcription of the original text can be found here: http://imaginaryinstruments.org/lovelace-analytical-engine/. Last accessed in 12 April 2023.

[2] For overview, see, for instance, du Sautoy (2019) and Miller (2019).

[3] For an overview, see https://iep.utm.edu/aesthetics-of-classical-music/#SH6c. Last accessed in 12 April 2023.

[4] For an analysis about the relationship between data sciences and AI, see, for instance, Kotu & Deshpande (2018).

[5] See, for instance: "[...] we have no power of introspection, but all knowledge of the internal world is derived by hypothetical reasoning from our knowledge of external facts" (PEIRCE, CP 5.265); "[...] we have no power of thinking without signs" (PEIRCE, CP 5.265); "[...] in my opinion it is much more true that the thoughts of a living writer are in any printed copy of his book than that they are in his brain" (PEIRCE, CP 7.364).

[6] More on this topic, see Atã & Queiroz (2019).

[7] The process here described is based on Schütz's analysis of the structure of making music together (SCHÜTZ, 1951). According to him, music meaning is a process of designing, communicating, and re-constructing an experience of inner time (*durée*).

[8] Here, I continue following the insights from Stjernfelt (2022: 70).

## referências bibliográficas


ADORNO, Theodor W. Current of Music: Elements of a Radio Theory. R. Hullot-Kentor (Ed). Polity Press, 2009.

ATÃ, Pedro, & QUEIROZ, João. O externalismo semiótico ativo de CS Peirce e a cantoria de viola como signo em ação. Trans/Form/Ação, v. 44, p. 177-204, 2021.

ATÃ, Pedro; QUEIROZ, João. Emergent Sign-Action. Classical Ballet As A Self-Organized And Temporally Distributed Semiotic Process. European journal of pragmatism and american Philosophy, v. 11, n. XI-2, 2019.

BELLUCCI, Francesco. Peirce's speculative grammar: Logic as semiotics. Routledge, 2017.

BONINI, Tiziano; GANDINI, Alessandro. "First week is editorial, second week is algorithmic": Platform gatekeepers and the platformization of music curation. Social Media+Society, v. 5, n. 4, p.1-11, 2019.

CHINOY, S. and MA, J. (2018). Why Songs of the Summer Sound the Same. Nytimes.com. Available in https://www.nytimes.com/interactive/2018/08/09/opinion/do-songs-of-the-summer-sound-the-same.html Access 29/05/2023.

COLAPIETRO, Vincent. The Music of Meaning. The American Journal of Semiotics, v. 36, n. 1/2, p. 11-45, 2020.

CUMMING, Naomi. The sonic self: Musical subjectivity and signification. Indiana University Press, 2000.

DE AGUIAR, Vinícius. Desenhando a música: análise da composição e da escuta segundo o conceito peirceano de diagrama. Tese de doutorado. Faculdade de Ciências da Universidade de Lisboa. Centro de Filosofia das Ciências da UL. Lisboa, 2022.

DIMOLITSAS, I., KANTARELIS, S., & FOUKA, A. SpotHitPy — A Study For ML-Based Song Hit Prediction Using Spotify. arXiv preprint, 2023. Available in https://arxiv.org/abs/2301.07978 Access 29/05/2023.

DU SAUTOY, Marcus. The Creativity Code: Art and Innovation in the Age of AI. Cambrdige: The Belknap Press of Harvard University Press, 2019.

FLUSSER, Vilém. Gestures. University of Minnesota Press, 2014.

FUBINI, Enrico. Estética Musical. Edições 70, 2008.

HANRAHAN, Nancy Weiss. Hearing the contradictions: Aesthetic experience, music and digitization. Cultural Sociology, v. 12, n. 3, p. 289-302, 2018.

HESMONDHALGH, David. Streaming's effects on music culture: Old anxieties and New simplifications. Cultural Sociology, v. 16, n. 1, p. 3-24, 2022.

KANT, Immanuel. Critique of Judgment. Trans.: W. S. Pluhar. Indianapolis/Cambridge: Hackett Publishing Company, 1987.

KIRCHER, Athanasius. Musurgia Universalis. Available in: https://imslp.org/wiki/Musurgia_Universalis_(Kircher,_Athanasius), 1640. Access in 29/05/2023.

KOTU, Vijay; DESHPANDE, Bala. Data science: concepts and practice. Morgan Kaufmann, 2018.

KRUSE, Felicia E. Is music a pure icon?. Transactions of the Charles S. Peirce Society, v. 43, p. 626-635, 2007.

MAASØ, Arnt; HAGEN, Anja Nylund. Metrics and decision-making in music streaming. Popular Communication, v. 18, n. 1, p. 18-31, 2020.

MANOVICH, Lev; ARIELLI, Emanuele. Artificial Aesthetics: A Critical Guide to AI, Media and Design. Available in http://manovich.net/index.php/projects/artificial-aesthetics. Access in 12/04/2023.

MANOVICH, Lev. AI Aesthetics. Moscow: Strelka Press, 2018a.

MANOVICH, Lev. 100 billion data rows per second: Media analytics in the early 21st century. International journal of communication, v. 12, p. 16, 2018b.

MAOR, Eli. Music by the numbers: From Pythagoras to Schoenberg. Princeton University Press, 2018.

MERSCH, Dieter. Epistemologies of aesthetics. Diaphanes, 2015.

MILLER, A. The Artist in the Machine: The World of AI-Powered Creativity. MIT, 2019.

MONELLE, Raymond. The sense of music: Semiotic essays. Princeton University Press, 2010.

NATTIEZ, Jean-Jacques. Music and discourse: Toward a semiology of music. Princeton University Press, 1990.

OLIVEIRA, Luís Felipe. A Emergência do Significado em Música. Tese de doutorado. Programa de Pós- Graduação em Música do Instituto de Artes da Universidade Estadual de Campinas – Área de concentração em Fundamentos Teóricos, Campinas, 2010.

OSBORNE, Richard; LAING, Dave (Ed.). Music by Numbers: The Use and Abuse of Statistics in the Music Industries. Intellect Books, 2020.

PAOLUCCI, Claudio. Cognitive semiotics: Integrating Signs, Minds, Meaning and Cognition . Springer International Publishing, 2021.

PEIRCE, Charles Sanders. Collected Papers of Charles Sanders Peirce, Vols. I-VIII, Charles Hartshorne, Paul Weiss, and Arthur Burks (Eds.). Harvard University Press (referred as CP), 1931-58.

PREY, Robert. Nothing personal: algorithmic individuation on music streaming platforms. Media, Culture & Society, v. 40, n. 7, p. 1086-1100, 2018.

SANTAELLA, Lucia. Nine modes of listening to music. Recherches sémiotiques, v. 37, n. 1, p. 51-60, 2017.

SCHEDL, Markus et al. Music recommendation systems: Techniques, use cases, and challenges. In: RICCI, Francesco, ROKACH, Lior, SHAPIRA, Bracha (Eds.). Recommender Systems Handbook. New York, NY: Springer US, p. 927-971, 2022.

SCHEDL, Markus et al. Music recommender systems. In: RICCI, Francesco; ROKACH, Lior; SHAPIRA, Bracha (Eds.). Recommender Systems Handbook. New York, NY: Springer US, p. 453-492, 2015.

SCHÜTZ, Alfred. Making music together: A study in social relationship. Social research, v.18, p. 76-97, 1951.

SERRÀ, Joan et al. Measuring the evolution of contemporary western popular music. Scientific reports, v. 2, n. 1, p. 1-6, 2012.

SHORT, Thomas Lloyd. Peirce's theory of signs. Cambridge University Press, 2007.

STJERNFELT, Frederik. Sheets, Diagrams, and Realism in Peirce. Walter de Gruyter GmbH & Co KG, 2022.

TARASTI, Eero. Semiotics of classical music. De Gruyter Mouton, 2012.

www.nytimes.com/interactive/2018/08/09/opinion/do-songsof-the-summer-sound-the-same.html. Access 29/05/2023.

ZALAMEA, Fernando. Peirce's Continuum: A Methodological and Mathematical Approach. 2001. Available in https://uberty.org/wp-content/uploads/2015/07/Zalamea-Peirces-Continuum.pdf Access 28/05/2023.