Alone Together. Some Reflections on Talking to ChatGPT, Listening to Music, and Private Language

Ondřej Beran
ondrej.beran(at)upce.cz

Abstract

This article draws on the author's experience of conversing with ChatGPT about music. These talks are often unsatisfactory. This observation does not aim to discredit the conversational skills of current advanced LLMs. Instead, the example of an elusive listener experience serves as an object of comparison, showing that conversations among human partners often feature borderline-private content (in the sense of Wittgenstein's private language arguments). Among humans, these gaps in meaning are bridged through structures of trust, reaching out to each other and various other "leaps of faith"; with ChatGPT, the private-language quality is exposed, as our conversations with chatbots do not rely on the underlying ground of trust, there is no friction and no sense of a productive conflict, and the chatbot cannot help us. This is a reflection on our (current) concepts, though, not a claim about the mindedness of present or future chatbots or predictions of what they will be able to do and when.

Keywords: ChatGPT, music, listener experience, private language, trust, conversation

Introduction

Advanced Large Language Models (LLMs), such as ChatGPT, are a hot topic nowadays. Naturally, for philosophical discussion, too. Although the widespread visibility of the phenomenon is still relatively recent (as of late 2022, when ChatGPT version 3 was launched), a varied philosophical response has already emerged. The existing pieces, I believe, capture at least some of the

most common (as far as I can judge) worries and concerns circulating in public and academic debates about the new wave of LLMs. They range from curious to serious and thorough.

Robert McGee published a flood of records of his conversations with ChatGPT on a remarkably diverse range of topics, including explorations of the use of ChatGPT in teaching philosophy, for instance, for designing reading lists or topic outlines for courses such as political philosophy (e.g. McGee 2023a, 2023b). There are a few more systematic discussions of opportunities (as well as threats) that ChatGPT represents for education, offered by Peters et al. (2024) or Tlili et al. (2023); cf. also the conversation between King and ChatGPT [sic!] (2023) specifically on the topic of plagiarism. Shortly after the public launch of ChatGPT, Luciano Floridi (2023) revisited the traditional question of whether we can talk about genuine intelligence in this context. Early critical voices pointing out that ChatGPT offers, at best, a semblance of intelligence, if not a "false promise", have come from various corners, ranging from academics (Chomsky et al. 2023, Hanna 2023) to writers (artists) (Hartenberger 2023, Sterling 2024). Recently, more comprehensive takes on the phenomenon appeared. Coeckelbergh and Gunkel (2024) argue that much of the (public and academic) debate about LLMs and the problem of authorship is driven by far too instrumentalist takes on technology and human language, and by a peculiar quasi-Platonic (and probably untenable) distinction between appearance and reality. Vallor's (2024) book-length discussion sidesteps the question of whether LLMs truly think by focusing on the fact that whether they think or not, they "mirror" powerfully (and often dangerously) the values of those who designed them, amplifying their racist, sexist, or capitalist biases not only in the produced texts themselves but also in the promoted range of ways they are used. There seems to be a degree of consensus that the gravest challenge posed by advanced AI may not be to the claim of the exceptionalism of the human mind and its capacities, but more straightforwardly political: corroding the popular political conversation and deliberation and trust in institutions (see Coeckelbergh 2024; he speaks of the corrosion of the "knowledge basis" of political deliberation) and twisting it to meet the economic interests of powerful players. The philosophy-of-mind issue of whether the thinking/intelligence claim is genuine does not have a straightforward relation to the AI's efficiency on the front of political life and social media; the success depends partly also on whether advanced AI is widely *understood and related to* by the broader public as genuine intelligence (or, soon,

AGI). This is one reason why an equal amount of genuine technological advancement and of a deliberate game of smoke and mirrors has been featured in the industry; cf. the earlier criticisms expressed by Crawford 2021, Parviainen and Coeckelbergh 2021 (or my own more neutrally phrased argument that the underlying ground of adopting a particular kind of attitude towards them needs to be built in order to make the practical ascription of intelligence possible, or at least smooth (Beran 2018; Beran 2022)).

Responses that rely specifically on Wittgenstein or Wittgensteinian philosophy are not that numerous (yet). Gubelmann (2023) suggests that while LLMs still fall short of linguistic understanding, there is no inherent obstacle to their eventual achievement of language mastery, as this consists in use. (Consider, on the other hand, Ground's (2024: 124n) cautious note about the possible need to distinguish "use" from "usage".) Often, writers in the Wittgensteinian tradition strive to employ, but also to (re)consider carefully the centrality of the human (form of life) in Wittgenstein. Abdel-Rahman (2024) relies on Wittgenstein to criticise the unjustified anthropocentric imagery of intelligence in Chomsky's and Brandom's criticism of AI. In contrast, Williams (2025) offers a nuanced argument for why relating to chatbots as persons represents only a derivative, "parasitic" case of an "attitude towards a soul". (In some respects, my remarks in the following sections are not far from points made by Williams.) Standish (2026) offers a complementary consideration of the (mis)application of the notions of "thinking" and "intelligence" to AI that may underappreciate the role of affectivity and that, in turn, may reshape our understanding of what "thinking" means even in humans.

My aim here is not to contribute to philosophical reflections on the nature and capacities of advanced chatbots. (Others can do and have done and will continue doing this much better than I.) I would rather exploit the newly emerged topic of conversing with advanced chatbots first to capture what is (at present, at least) different from conversations with humans and then to illuminate a particular area of human conversations that is perhaps not discussed as centrally by the philosophy of language as it could be, yet is – I believe – related to what is at the heart of human interactions. (Among the current philosophical writings about LLMs, an intriguing parallel to my object of interest is offered by the discussion of the limits of using AI in education, by Jin and Wang (2025), who argue that what is lacking in interactions with AI

is the quality brought about by two interacting *subjects*, along with all that it involves.)

My central, albeit anecdotal, example, introduced in Section 1, is the series of my attempts to talk to ChatGPT[1] about music and, specifically, about listening tips for a particular kind (mood?) of music. The hilarious and unsatisfactory results may suggest that I have been chasing for contents of a "private" kind in a Wittgensteinian sense. After briefly introducing Wittgenstein's private language argument (in section 2), I suggest that a relatively significant part of our conversations about music (listener experience) tends to move in private, borderline-private, or private-prone territories. I then suggest that while in conversations with chatbots, this private character tends to be laid bare, if not exacerbated (section 3), in conversations with people (friends), it can be fruitfully transformed into enriching conversational experiences (section 4). I suggest that the setting of trust, which underlies human conversations, is at play here. In section 5, I conclude by offering a few cursory notes about the nature of human speaking to each other in general, building on the post-Wittgensteinian notion of conversation and suggesting that while we may fail to have genuine conversations with artificial entities, this is not necessarily due to prerequisites missing on the part of chatbots (I do not offer any hypotheses about that), but it has to do with attitudes we adopt (or can adopt, for now) towards them and with the difference in the relational contexts in which the interactions with other human and with LLMs, respectively, take place.

1. Example: Emotions Expressed by Music

Like many people today, I occasionally use ChatGPT (first the version 3.5, then 4, then 4o). I find many of those chats useful, occasionally even refreshing; for instance, when it explained succinctly to me (and to my daughter, the actual originator of this query) all the purposes for which octopuses use the suckers on their tentacles. The chatbot often displays a remarkable restraint, for instance, when I asked it to suggest a viable business model for a vegan-friendly modern-day version of haruspicy. (Admittedly, the task set was borderline

[1] To a smaller extent, I replicated some of the below-described experiments with Google's Bard and Claude 3.5 Sonnet, with similar results. For the sake of simplicity, I will keep talking only about ChatGPT; this simplification should not pose a problem, as my aim is *not* to assess the real capabilities of particular advanced LLMs.

absurd, and the chatbot responded correspondingly – perhaps also because it is trained to avoid contested territories such as religion or sensitive issues such as animal rights.) Other times, when I ask problematic questions, such as those regarding politics and economics, the chatbot tends to reply in a non-committal, "balanced" manner. I could not make it assess the Pinochet regime in Chile in a way that would *not* list both its pros and its cons.

I have also made several attempts to consult ChatGPT as a "music curator": for instance, I have given it the names of a few of my favourite bands and asked it to suggest "something similar" to me. I have been exploring this inquiry for some time, on multiple fronts. In a separate thread, I asked ChatGPT to suggest bands or songs that would exemplify "melancholic euphoria," a term I coined to capture the quality I was looking for in the "something similar" in the initial query. In another thread, I tried to make ChatGPT give me particular examples of songs by my favourite bands displaying (in my view) this mixed mood clearly, and then I asked it to suggest songs by different artists in a similar mood. Although I – patiently, by my standards; well, I was patient also because I was driven – rephrased my inquiry several times and tried to develop it into a form more focused on what I was up to, the results remained confusing and unconvincing to me. I almost never managed to squeeze out of it a suggestion that would, in a human interlocutor, make me feel they have a hunch of what I am up to, though I was speaking to an entity handling the words used with apparent competence.

Mostly, ChatGPT offered support similar to cases when I cannot find/remember a word/a name and somebody suggests, "Let us list the most probable names circulating in this context: this, and this, and this, … – which one would you say it was?" (or "let us try the alphabet and see whether we can find what the first letter was – was it A…? Or B…? Or C…? …") I am not saying these methods cannot help finding a temporarily lost word, for instance, the momentarily not recalled name of a person who nevertheless exists "out there". However, importantly, the name of an actual person can, with all likelihood, be found by more than one method, which helps narrow the circle. Finding a forgotten word from a *dream* in this way would prove much more difficult since words from a dream are locked against this kind of approach much more powerfully than temporarily forgotten names of actual people. There are no alternative means to support the access. After a while, you no longer feel sure whether you really dreamt a dream in which this was

happening, partly because notions like "remembering," "happening", or "sure" do not apply to dreams as neatly as they do to most areas of our (awake) life.

Also, trying to capture a word from a dream often comprises both the aspect of identifying it and the aspect of understanding what it is about the word that drives your endeavour to identify it. Why it was important in the dream and why it is important for you now. The above suggestions are of little help in both these respects. In the former, because dreams are well (though not bullet-proof) protected against these techniques. In the latter, because the list of the most probable words featuring in dreams is strangely uninformative regarding what one wants to find.

This parallel with dreams is not without importance. Imagine you tell ChatGPT something like "Pack Yr Romantic Mind by Stereolab is a good example of 'melancholic euphoria' in a song; could you explain this quality to me and suggest songs that would work similarly?" (This is a shortened and simplified version of one of my prompts.) ChatGPT duly produces a detailed description of what makes the song particularly melancholic-euphoric, which is, however, rather uninformative: telling you that "it combines upbeat elements with haunting undertones", dropping words like "dreamy" or "ethereal". And its suggestions of similar songs, mostly, just seem to be missing the point: for instance, my older conversations with ChatGPT 3.5 yielded random-sounding suggestions like "Karma Police" by Radiohead ("Yeah, great song, but… eh") or "Rebellion (Lies)" by Arcade Fire (just "…eh"). When I spoke with Claude 3:5 Sonnet more recently, it mostly avoided blatantly bizarre tips, but it was like conversing with a skilled user of Wikipedia entries about shoegaze or dream pop.[2] Consider this parallel: somebody volunteers to tell you a dream story presumably analogous to a particular dream you have had but cannot now clearly remember, in order to help you remember, but they will rely only on a rather general, textbook knowledge of what dreams are. My later conversations with ChatGPT 4o mostly did not feature any blatantly "meh" tips; in a few isolated cases, it even suggested songs I already had on my playlist (but did not mention as examples in the prompt). However, it did not expand my horizons in any way – the least misplaced suggestions were things known to me, and the suggestions unknown to me were typically far further

[2] As I was using Stereolab as a reference point, it is worth noting that my AI interlocutors (I conversed with them in English, which may have been the explanation) wholly disregarded the partly French roots of Stereolab and only suggested Anglo-American – or Nordic – artists, ignoring the whole continent of French avant pop (La Féline, April March).

from the mark. What did *not* happen was a discovery of something unknown that would, so to speak, blow my mind, whether in meeting the initial request spot-on or despite *not meeting* it.[3]

This example is not designed to blame ChatGPT for its ineptitude in conversing about music with insight, for there was something unfair about the inquiry. What interests me is how lost I felt face-to-face with such feedback or in such a dialogue. In fact, I felt slightly stupid: What was I expecting? However, this is partly due to the topic's peculiar character, which really has something in common with our dreams – contents displaying qualities that nobody else can access independently and comparably.

Thus, one more (one last) recourse to my conversations with ChatGPT, this time about perception: when I asked it, in another, unrelated thread, to provide me with a few examples of real-life (actual) experiences that are based on equally precarious foundations as "private language", as Wittgenstein famously discusses it (PI: §§ 201f, 243ff), it offered me, among other things, this (now I am quoting *verbatim*):

> Private symbolic associations: Someone might create a language where words or symbols are assigned based on personal and idiosyncratic associations. For instance, they might assign the word "tree" to represent the sensation of warmth.

One might argue that my dogged attempts to use the words "melancholic euphoria" to refer to a particular kind of listener's experience that I am unable to expand in any way, and that the entity I talk (or "talk") with about it seems to be missing constitutionally, are of precisely this kind: based on a private symbolic association. Is that so?

2. The Private Language Argument

The most straightforward entry point to Wittgenstein's private language argument is the famous, relatively simple imagined scenario from § 258 of the *Investigations*: of a person keeping a "sensation diary" in which they record every appearance of a particular sensation S (by writing down "S"), as it appears to their mind's eye, so to speak. It is essential, though, that Wittgenstein is not speaking about psychological experiences in the standard sense, as when we

[3] The one musical discovery that I made, thanks to conversations with ChatGPT, and for which I am truly grateful – Barbara Morgenstern –, emerged in an unrelated thread. In a manner intriguing for me, her music would both fit and playfully evade the melancholic euphoria request; for once, the narrow horizons of my mind have been broadened.

routinely talk with other people about anger, anxiety, or joy that we experience, or about our sensations of heat or bright blue we have (when we sit at the seaside in summer). The definition of "S" cannot be formulated, and S cannot be pointed to in a usual sense. Consequently, we have no reason to call S a "sensation," for "'sensation' is a word of our common language, not of one intelligible to me alone" (PI: § 261). In fact, we do not even have the right to say of S that we have a "something" because "(having a) something" is also an expression belonging to our common language in the same way as "feeling" or "sensation". After all, "a nothing would serve just as well as a something about which nothing could be said" (PI: § 304).

The exact point of the private language argument has been subject to endless debates (see the overview in Candlish 2019 or the recent subtle and detailed reconstruction of the Wittgenstein-Rhees position in Munz 2024). Some voices in Wittgenstein scholarship have suggested that Wittgenstein aims to argue that such a private language, as his example outlines it, is impossible (e.g. Gert 1986, Sussman 1995). However, it rather seems that Wittgenstein wants to point out that our language, as we know it and speak it – including those of its territories that concern "inner" sensations and experiences – cannot have its meaning based on such foundations as the example of § 258 of the Investigations entertains. The foundations of meaning must be public, intersubjective; otherwise, there would be no "criterion of correctness", and one could not make a mistake because whatever seems right to them would be right, and "that only means that here we can't talk about 'right'" (to quote PI: § 258). While this reading certainly does not exhaust the richness of the subject and of the following debate, and some might find it too simplistic, it is not controversial, and many voices in the scholarship can be understood as representatives of versions of it, ranging as widely as, for instance, from Rhees (1954) to Baker and Hacker (1984).

Wittgenstein himself considers various kinds of private language, including the Ayerian examples of the solitary speech of Robinson-like characters. However, most of these have not made it from the MSs to the text of the *Investigations* (see Munz 2024: 243ff, esp. p. 250f). There, he primarily focuses on private language as a peculiar side hobby of competent speakers of public language, which, *by definition*, another person cannot understand. In fact, the very idea of coining such a language is described by Wittgenstein as an unfortunate variation (extrapolation) on normal, public language, somewhat in

the manner of some advances in formal logic.[4] The way in which our linguistic interactions with LLMs both are and are not like interactions with other people perhaps represents another case of this extrapolation. (I will try to flesh out this suggestion a bit more in the following sections.)

The discussion about the private language argument(s) evolves into many fine ramifications, into which I will not go, though, for my point here is not the exegesis of the argument as such. My aim is to consider which aspects of Wittgenstein's discussion of private language(s) shed some light on the phenomenon of talking to LLMs. In general, the problem of the non-functionality of a "language" based on private "criteria," identified as the core of Wittgenstein's argument in the *Investigations* by Baker and Hacker, Rhees, or Munz, is close to the core of my interest in conversations with LLMs (see the following section). However, I would also like to briefly revisit the formerly mentioned point, which is occasionally (though probably mistakenly) attributed to Wittgenstein: that private language, as outlined in the example of the private sensation diarist, is impossible. It is quite clear that anybody could amuse themselves with such an exercise if they want; the only problem (the real problem) is that the results cannot be plausibly called "language". Yet, there are, in fact, not only such fictional examples but various territories of existing language practice that resemble Wittgenstein's example and that surround the ocean of properly working language games as fringe, marginal phenomena or occasionally hide in its centre as peculiar blind spots – for instance, glossolalia (Samarin 1972), "twin talk" (Zazzo 2009), "schizophrenese" (Wolcott 1972), or more mundane situations such as that of somebody's feeling that they are constantly and notoriously misunderstood by others given what they "wanted to say" (I discuss these examples more systematically in Beran 2012). For sure, the vast majority of our language practices, or their core, if there is any, cannot be like that; but still, how we use language is far from unified. Instead, it proves porous, heterogeneous, and features areas that are distinctly "fishier" than others.

I have, of course, not proclaimed "melancholic euphoria" a sign the definition of which "*cannot* be formulated". However, I am not entirely sure that it *can* be formulated either. My repeated attempts to specify my prompts, as well as the feedback that I was getting from ChatGPT, strongly resembled groping for outward criteria of privately construed inner experiences. And my

[4] I am indebted to one of my reviewers for drawing my attention to this parallel.

sense of dissatisfaction opened room for the suspicion that I might not be satisfied with *any* description of this kind because anything provided by anybody (or any*thing*) other than myself would be inherently ill-fitted to the topic. Might this make *my notion* of "melancholic euphoria" a notion parasitising on genuine psychological concepts in a "private" manner, therefore, an empty non-notion? Partly, I am afraid, it is the case. But not wholly, and a part of the reason may have to do with the nature of our "conversations" about such topics with chatbots as opposed to conversations with other people. More in the following two sections.

3. Talking (Music) with ChatGPT

In the previous sections, I mentioned aspects of my conversations with ChatGPT about music that resembled a private language situation. I also suggested that it had to do with the nature of the talked-about subject. Yet, this is not all. My dissatisfaction with ChatGPT answers also echoes some of the differences that Theodor Adorno finds between types of music listeners (1976: chapter 1). ChatGPT resembles what Adorno calls a "culture consumer": somebody whose relation to music consists mainly in their access to extensive collections of records and a quantity of information *about* music, rather firmly anchored in the widely received assessments of the discussed pieces. Hence, perhaps, my dissatisfaction; it might have been certainly less bitter if the chatbot had talked to me in a way that would strike me as the kind of feedback one gets from those whom Adorno calls "good listeners" – people capable of spontaneously and perceptively understanding and judging music also concerning the meaning it carries and develops through the employed modes of composition. (An understanding that is internal/qualitative rather than external/quantitative.) Still, even that hope may have been deluded; perhaps I would have reacted to *any* offered comment with a private-based dissatisfaction. At any rate, ChatGPT reacted in a way that was not quite at the level I would have wanted it to be. Yet, importantly, this desired level is *not* fictional, at least in some cases of talking about music.

More notably, though, ChatGPT has failed significantly in the second part of the inquiry: in its musical recommendations. I asked it to broaden my horizons in a way that would make me happy, which is certainly not an impossible goal. When we host music-playing parties with friends or attend music concerts with them (based on their suggestions), this is what happens all

the time. We are being suggested musical tips previously unknown to us, and we fall for (some of) them. And sometimes, it's music like virtually nothing we were fond of before. Other times, these experiences build upon our previous listening experiences. They show us things in important aspects similar to what we already appreciate, but also such that refine our tastes further, even in unexpected directions. Honestly, our conversations about music with friends can also be unsatisfactory and lacking in analytical insights into the characteristic aspects of the music that inspires us. However, oddly enough, we continue to discover new things thanks to these conversations, even so.

Why doesn't this work with ChatGPT? (Or why was it not working for me?) A simple answer would be: Conversations with ChatGPT represent a context in which nobody helps or pushes you to anchor your borderline-private comments in a way that would allow them to work as non-private. Using a term taken from the public language (such as "melancholic euphoria") to refer to a psychological experience that you remain uncertain about – uncertain whether it can be "defined" or "pointed to" at all – leaves you in an unsteady position. Asking ChatGPT about "melancholic euphoria (in music)" will likely lead to it giving you answers similar to what you would find in a dictionary, which will duly define what "melancholic" and "euphoria" mean and what their combination may therefore mean (perhaps using "moll" or other musical descriptors and their explanations). ChatGPT will ingenuously recombine these partial inputs into what will look (and work) like a single integrated original answer. It will spare you the work of browsing through the pages of these volumes, making excerpts, and arranging them together. However, to the extent that any such commentary strikes you as – regarding the topic – inherently misplaced or vacuous, this places you, practically, in a position analogous, with regards to your "melancholic euphoria", to that of Wittgenstein's private sensations diarist. The help you think you require here needs to come from somebody who has (or knows what it's like to have) the same kind of access as you have, to the same kind of contents.

That, crucially, does *not* need to be a problem; I will try to indicate how in the following section. Here, I want to suggest that it is a problem with ChatGPT because whatever it does (says) does not assume the meaning of the "help coming from somebody who knows what it's like…" For one thing, nobody urges you here to be more specific or clearer, glosses over the gap, or starts making fun of your unclear rambling about melancholic euphoria (or

commiserates with you thanks to a shared experience of once having had such a moment of "insight" and lost it again, too). It is not that ChatGPT could not utter what could be read as a request for you to be more intelligible about the nature of the relationship between the words "melancholic euphoria" and the (inner) listener experience that you claim you have. It can do that, though you would have to ask it to talk to you "like that" (or a model trained to speak in that way by default would need to be launched). But only when this kind of talk comes from somebody we can understand as the "somebody who knows what it's like…", it truly means urging, glossing over, or commiserating. Urging requires further pieces of setting that are simply absent here. Even if ChatGPT duly urges you (as prompted) to be more specific, it does not make any difference to your being left with the nagging worry that "whatever is going to seem right (or wrong) to you is right (or wrong)". The external prop I am creating (prompting) here for myself resembles the examples, mocked by Wittgenstein (PI: § 265), of correcting (justifying) a memory by another memory or confirming the veracity of a newspaper report by buying several copies of the paper.

Urging is a demand we either acknowledge or disregard, and when we disregard it, we respond with annoyance, amusement, or stress (for instance). We may still lack the conceptual and imaginative resources required to relate to our conversations with chatbots in terms of being thus urged by them, and we still don't know what it would mean to say that chatbots themselves have (or do not have) these resources to relate to us correspondingly – to respond with the same kind of variety to our urging.

A related part of the problem is that dictionary-like descriptions are ill-suited for helping you in a situation where you want to understand and unpack a private-like experience. These descriptions do nothing to dispel your feeling that something else would have to happen so that it is not only you who can understand the experience, or so that you don't end up feeling that not even you can be sure you can access the experience. (Remember the case of the failure to recall the word from a dream.) You could do the same – checking a dictionary – on your own and still be left where you were before. Not because you cannot try to analyse, describe and label your psychological experiences *on your own* (some come with *In Search of Lost Time* as the fruit of such endeavour), but because checking a dictionary is *not* the way to do this.

For these reasons, one tends to end up in a private language situation here. Not because ChatGPT is *by its nature* no real interlocutor (or even because it is thereby *proved* to be no real interlocutor), but simply because what one does in a conversation with ChatGPT about an elusive subject does not differ from being alone in the echo chamber of one's thoughts, indistinguishable in this context from what one thinks are her thoughts, and having only dictionaries and encyclopedias to rely on.

4. Talking (Music) with Friends

What is different when I talk about melancholic euphoria in music with a friend? As I suggested earlier, neither of us, both presumably human, might be able to produce a very good analytic insight into the notion. However, we may be both driven, even excited, by the discussion, and the "objectively unsatisfying" quality of our comments offered so far (as assessed by someone who would read a transcript of our debate) does nothing to discourage us from throwing in and trying further listening suggestions.

Notably, we are still moving on "private" ground; not even here has the notion of "melancholic euphoria" established itself as something clear, solid, definable, and pointable (yet). This may not change throughout the conversation. Yet, instead of the eventual demoralisation of discussing musical tips with ChatGPT, this is, in fact, fun, and on occasions, unexpected tips of music fitting the searched-for characterisation and enthusiastically agreed upon as such by both parties keep being found. We may not be able to explain this fortunate encounter; I feel the need to stress that the featured kind of music talk here is lay and, in important respects, uneducated. In his inquiry into listening to music, Martyn Evans (1990: esp. chap. 2) stresses that a listener's emotions arise from understanding (p. 45), which we routinely unpack – indeed, should be able to unpack – through verbal judgments. Yet, Evans concedes that it is difficult to identify where even such generic emotional characteristics of music as "sad" or "joyful" are located, and he ultimately roots them in the soil of shared patterns and forms of understanding, that is, of life, with music. But where this rootedness is of little help – in cases where a characterisation (felt as fitting) of a piece of music is of an elusive nature (the "melancholic euphoria"), and without much certainty as to further instances of this kind – one is in a different situation. Unlike Evans' examples (p. 49), my

example is one of fumbling clumsily for words but also, crucially, of enjoying and embracing the fumbling.

The example of fumbling clumsily for words and explanations while "we are, after all, among friends" suggests that something very different is happening with borderline-private features of language use in a conversation with another *person*. With our fumbling comments, we try to reach out to the other; or rather, we act in good faith that what we say is intelligible as trying to reach out to each other. And we respond in the same way to what others say to us. Listening tips from a friend may not strike me as better than those offered by ChatGPT. If my friend suggests listening to "Rebellion (Lies)" by Arcade Fire to me during a "melancholic euphoria" music party, I might respond by either laughing it off ("Now that's the most bonkers thing you have said tonight! Why not Ricky Skaggs[5] straightaway?") or by getting over it, as their having a "weaker moment". In fact, that some of the suggestions that we both propound turn out to be hilariously misplaced is a part of the fun; this happens in a conversation among people you feel comfortable with when you explore and enjoy together the absurdity, indeed, these encounters quite fruitfully help us build our (new) friendships.[6] On the other hand, when ChatGPT offers what would strike me, from a human interlocutor, as hilariously misplaced musical tips, I first wave it off and then stop trying because it's becoming tedious. I do not try to reach out to it because I do not know what it would mean here. (Friendship is not engendered here.) With a human interlocutor, I can be amused by their not seeing, at first, what is so amusing here, and then I can show them. Or, I can be angry or feel let down if they do not "give" the conversation as much as I do and all such efforts of mine fall flat on their ears. I do not know what it would be like to be analogously let down by ChatGPT or to *show* it that its tips are amusing.

Conversely, when the conversation does go well – when my friend and I keep trying despite grappling with a borderline-private characterisation and when we enjoy it – that is not because a human interlocutor constitutionally has a mental setup better suited to help me out of a private trap than an artificial

[5] No offence to Ricky Skaggs intended. Years ago, when a friend of mine and I happened to hear Ricky Skaggs' performance at a music festival, after the artist ended, my friend (a person with perfect pitch) said, "That set did not have a single moll tone".

[6] This observation is partly borrowed from a nice point made succinctly by one of my anonymous reviewers.

interlocutor. This is because our reaching out to each other reflects that the setting of the situation differs fundamentally.

Some authors explain this as an aspect of *trust*, which is the foundational setting of how we relate to each other, of our conversations (see Cockburn 2014, Lagerspetz 2015). Trust, as Olli Lagerspetz observes, is not an experienced psychological state (a feeling) but rather an aspect of the whole interconnected web of our attitudes towards each other (2015: 96). And even though contents such as "melancholic euphoria" may behave as "private," or borderline-private, that needn't be such an obstacle, to the extent that we realise that whether the borderline-private nature of certain topics turns out to be an inescapable trap is not necessarily the matter of their intrinsic qualities, but rather of how we can (or cannot) commit to a particular manner of conversation about them (cf. Lagerspetz 2015: 119). With other people, and especially with people we *particularly* trust, we feel safe to conduct "leaps of faith" of maintaining a conversation about topics about which we are not able – at first, but perhaps not only at first – to tell much more than Wittgenstein's private diarist can tell "S". And the maintained conversation often eventually bears fruits, which, from the viewpoint of a putative "private language argument purist," would look inexplicable; certainly unexpected. However, from another point of view, questionable as the conversation topic is, it serves as a foundation for very real enthusiastic conversational encounters over the discussed qualities and topics, elusive as they may be.

Discussing the stubbornly elusive quality of melancholic euphoria in music with a friend can turn out to be a situation like this. In a sense, though I cannot quite explain why this conversation works – or even *whether* it works in a manner comparable to exchanges of pieces of information and whether it *really* promotes and develops our intuitions about this kind of music – it is not such a big deal because the point of talking to each other is not exhausted by whether we "objectively get something" out of the conversation. Oddly enough, whether the meaning of "melancholic euphoria" can be established intersubjectively – as Wittgenstein's private language argument (in Baker and Hacker's reading) stresses it should be – is not the point here. As David Cockburn (2014: 57) points out,

> Talking with another is not, or not solely, a means to something else. Conversation is a form of contact between people; and in so far as I have no interest in the *person*

> – in so far as my interest in what she says, in what she *means* by her words, relates purely to my need for information or my ability to predict or control her future behaviour – there is a sense in which I am not talking with (talking *with*) her at all.

From this point of view, ChatGPT's incapacity to tell me much that would strike me as an illuminating insight into melancholic euphoria in music is something that matters (unfavourably) to me much more than my friend's incapacity or my and my friend's joint incapacity even after a long debate and a long list of records played.

Fundamentally, the same kind of "undetermined-yet" conversation contents is approached among friends as, precisely, undetermined-*yet*, as a temporary shortcoming in an otherwise sound interpersonal relationship of making ourselves intelligible to each other. The indeterminacy is often never really thematised or focused on because it may not be the point of their conversation. On the other hand, with a chatbot, these undetermined contents are immediately approached with scepticism, with the view of their indeterminacy as i) centrally salient and characteristically marking the interaction and ii) quite likely inevitable and undesirable. In one case, trust is implicitly taken as a given (even if it temporarily or for a longer while bears no "fruits"), as reflected by the interconnected manner in which we act, feel, perceive, and respond to each other. In the other case, it is explicitly taken as a challenge to conquer, accompanied by a foundational uncertainty about the prospects of the endeavour. (Note that in this respect, the typical Turing Test settings differ from encounters among humans; perhaps the presence/absence of *relatable* embodiment – "the best picture of the human soul" – on *both* sides is what makes the difference.) In the latter case, we take trust to be something that needs to be rationally justified; in the former case, we do not, because trust between human beings does not have to fulfil a purpose (Lagerspetz 2015: 155), just as conversation needs not (Cockburn 2014: section 4).

I cannot quite put my finger on why these aspects are missing from our conversations (or "conversations") with ChatGPT. But they are. I will offer a few suggestions in the following final section.

5. Talking, with ChatGPT and Otherwise

In the opening of his "Art and Philosophy," Rush Rhees warns against the idea (not uncommonly attributed to Wittgenstein) of philosophy as a clarification

of the (mis)use of words (1969: 133ff). As he argues, if a philosopher wants to address confusion regarding, say, "time," the thing to do is *not* to consult a dictionary (the dictionary entry "time"): "[P]hilosophical difficulties are rooted not so much in 'linguistic confusions' as in confusions about *language*"; the problem being not "how words are used, but what using words is" (p. 134). In the case of Rhees' example, "What time is" is a question about the role time plays in the lives of people: beings who speak, also about time, among many other things.

Analogously, in our case, if "melancholic euphoria" is a thing interesting (viable) to talk about, then the fruitfully contested ground does not invite dictionary-like explanations of the label but rather lies in explorations of what role is played in our lives by the kind of situations such as that in which the music-loving friends from our recurring example try to clarify or solidify their relation to a particular thrilling aspect of some music using the term "melancholic euphoria". (Notably, we can learn much that is interesting and relevant about this aspect even when reflecting on cases where the term does not feature.) In this sense, "melancholic euphoria" is also something that can be shared with another person and it is melancholic and euphoric, exactly, due to this shared experience, recognised as shared and as something made possible through this (recognised) sharing. Elusive but elating ways – defying a fixed description, explanation, and analysis – of experiencing music, whether shared or solitary, are often occasioned by various kinds of mind-altering settings. The most mundane would be so prosaic as fooling with a pub jukebox after a few beers; but equally common and typical, only much more complex, would be the trance-like experience of being together at a music show, dancing, listening and singing, perhaps also drinking, smoking pot, or taking MDMA. Some authors find a link between these states of mind-alteration by a substance and/or arts on the one hand and *bona fide* spiritual experiences on the other hand; be it with a critical spin (Ross 2013: 42f) or with an open-minded appreciation for states of under-the-influence elation, as, for instance, in the classic observation by William James (1985: 397).

This is not to denigrate what chatbots can do by pointing out the obvious (but in itself, not that interesting) fact that a human being and a chatbot cannot get drunk or high together at a music show. (Yet.) This is rather to point out that human language interactions can sustain much more precariousness, open-endedness, imbalance and dynamics than Wittgenstein is sometimes taken to

admit, and certainly more than our interactions with LLMs. A part of what makes the human interactions (feel) alive and thrilling is that there is a whole network of paralinguistic contexts, actions and interactions that also allow borderline-private areas of our language to be surprisingly meaningful and productive. This network is, at least for now, exclusively a human domain. (It can and does include interactions with animals, though.) In the Investigations, Wittgenstein argues that we can "bridge" the episodes of a lack (or breakdown) of sense thanks to the "reference point" provided by the "shared human behaviour" (PI: § 206). This shared ground also allows us to make a good (and varied) sense of the mistakes we have: the mistakes are mistakes not because of intrinsic properties of the particular linguistic moves (particular silly listening suggestions), but because they exist as such within a larger context of "customs, habits, institutions" (PI: § 199). This larger context offers avenues to pursue some of the mistakes we make as trivial, some as corrigible, some as eye-opening, some as cute and exhilarating. This context provides a point of reference for our mistakes that makes it possible for our friends' silly listening tips to be silly listening tips, while the exact same words written by chatbots remain in a vacuum. Pelland et al. (2024: 191ff) offer an interesting systematic explanation to this phenomenon: the missing shared ground prevents the kind of mistakes committed by LLMs from being interpretable as systematic and makes them essentially "ruleless" (regellos) (Pelland and his co-authors refer to this distinction made by Wittgenstein in PI: § 143).

The points I have sketched in this text, thus, besides having only a very limited general reach, were not designed as an observation about what LLMs can or cannot do or will be able to do. (Or even whether we should be worried about it.) I do not know this; I only expect us to keep being surprised at a surprisingly fast pace. I only suggest that the experiences of talking with ChatGPT can help us see more clearly what is distinctive about talking with other people and, by a philosophical extension, perhaps about speaking in general.

It has been observed several times that ChatGPT tells you what it has learnt (or been trained) to guess as that which will "meet" or "satisfy" (in whatever way) the prompt, and that it does not shy away from "hallucinating" plausibly-looking nonsense (cf. interesting criticisms by Li 2023 or Alkaissi and McFarlane 2023). To be fair, many human conversation partners behave similarly. However, poking and probing each other regarding whether they can

plausibly elaborate on their initial replies to task assignments (assignments that require providing factual information) is not a representative picture of human linguistic interactions. Not necessarily because it's non-representative statistically – I have no idea – but because there seem to be good philosophical reasons to assume that the foundational level of human (linguistic) interaction lies elsewhere. As some post-Wittgensteinian thinkers mentioned here (representatively, Rhees 2006) – but also, for example, Gadamer (1975: 443) – argue, language consists, deep down, in *conversation*. And conversation is *not* there for a particular purpose; we do not speak to reach particular aims, but we are reaching particular aims as making ourselves mutually intelligible in speaking language makes this (along with the very coining of such an idea) possible for us.

Underlying conversation is the attitude of trust, allowing us to reach out to each other and try to go along, even when we do not quite know what we are talking about for a while. This is what happens – what *can* happen – when contents behaving in a private-like way are introduced or alluded to in conversation. In conversation with ChatGPT, the private aspect is haloed and laid bare, and the conversation easily ends up with one's feeling like being locked within a private language. The reason is that there is no reaching out to each other or holding each other accountable here yet; no trust binds us together here with the putative "other". There are no "leaps of faith" here to be performed *together* or *towards another*. Whatever the actual capacities of the currently developed LLMs are, it is *we* who still seem to be lacking the necessary conceptual and imaginative resources for adopting an "attitude towards a soul" towards them (cf. Beran 2018, Beran 2022, Danziger 2022, or Hertzberg 2020). Or, in other words, while trust with humans represents a given that can be, episodically, exploited rather forcefully or even eroded and breached, there is no breach of trust committed by "not understanding" chatbots, because there was no underlying trust from the very beginning. (It needs to be built; and in this respect, it seems foolish to expect our interactions with AI ever to mirror, 100%, human-human interactions. But it *doesn't* mean that, *if* trust can be built, human-AI interactions could not, in important respects, resemble the richness and complexity of human interactions.)

I thus offer these reflections here primarily as observations about the kind of language games that are unsteady and pervaded with borderline private conversation stakes. ChatGPT behaves (or we behave in our interactions with

ChatGPT) in these situations in a way that uncovers and exposes the precarious nature of these interactions, but does nothing to bridge the breach. In contrast, such interactions between humans are open-ended, betting on the chance that against all odds, the private pieces can click into their place and "words for them will emerge like flowers" (to echo Heidegger's discussions of Hölderlin); which indeed happens, surprisingly often.[7]

Bibliography

Abdel-Rahman, L. 2024. "The Forms of Artificially Intelligent Life: Brandom, Chomsky and Wittgenstein on the Possibility of Strong-AI". In B. Ball, A. C. Helliwell, and A. Rossi, eds., *Wittgenstein and Artificial Intelligence, Volume I: Mind and Language*, London: Anthem Press, 105–121.

Adorno, T. W. 1976. *Introduction to the Sociology of Music*. Translated by E. B. Ashton. New York: The Seabury Press.

Alkaissi, H., and S. I. McFarlane. 2023. "Artificial Hallucinations in ChatGPT: Implications in Scientific Writing". *Cureus* 15 (2), e35179, https://doi.org/10.7759/cureus.35179.

Baker, G. P., and P. M. S. Hacker. 1984. *Scepticism, Rules and Language*. Oxford: Blackwell.

Beran, O. 2012. *Soukromé jazyky*. Praha: Filosofia.

Beran, O. 2018. "An Attitude Towards an Artificial Soul? Responses to the 'Nazi Chatbot'". *Philosophical Investigations* 41, 42–69.

Beran, O. 2022. "Conversational Machinations". In: S. Aldrin Salskov, O. Beran, and N. Hämäläinen, eds., *Ethical Inquiries after Wittgenstein*, Cham: Springer, 2022, 255–272.

Candlish, S. 2019. "Private Language". In: E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/entries/private-language/

Chomsky, N., I. Roberts, and J. Watumull. 2023. "The False Promise of ChatGPT". *The New York Times*, March 8. https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html.

Cockburn, D. 2014. "Trust in Conversation". *Nordic Wittgenstein Review* 3 (1), 47–67.

Coeckelbergh, M. 2024. *Why AI Undermines Democracy and What To Do About It*. Cambridge: Polity Press.

Coeckelbergh, M, and D. J. Gunkel. 2024. "ChatGPT: Deconstructing the Debate and Moving It Forward". *AI & Society* 39, 2221–2231.

[7] This work was supported by the European Regional Development Fund project "Beyond Security: Role of Conflict in Resilience-Building" (reg. no.: CZ.02.01.01/00/22_008/0004595).
I would also like to thank my reviewers for their useful comments and suggestions, not least for the listening suggestions they have generously provided.

Crawford, K. 2021. *Atlas of AI*. New Haven: Yale University Press.
Danziger, S. 2022. "Intelligence as a Social Concept: A Socio-Technological Interpretation of the Turing Test". *Philosophy & Technology* 35 (3), article 68, https://doi.org/10.1007/s13347-022-00561-z.
Evans, M. 1990. *Listening to Music.* Basingstoke: The Macmillan Press.
Floridi, L. 2023. "AI as Agency Without Intelligence: On ChatGPT, Large Language Models, and Other Generative Models". *Philosophy & Technology* 36 (1), article 15, https://doi.org/10.1007/s13347-023-00621-y.
Gadamer, H.-G. 1975. *Truth and Method.* Translated by Joel Weinsheimer and Donald G. Marshall. London: Continuum.
Gert, B. 1986. "Wittgenstein's Private Language Arguments". *Synthese* 68, 409–439.
Ground, I. 2024. "Black Boxes, Beetles and Beasts". In B. Ball, A. C. Helliwell, and A. Rossi, eds., *Wittgenstein and Artificial Intelligence, Volume I: Mind and Language*, London: Anthem Press, 123–144.
Gubelmann, R. 2023. "A Loosely Wittgensteinian Conception of the Linguistic Understanding of Large Language Models like BERT, GPT-3, and ChatGPT". *Grazer Philosophische Studien* 99, 485–523.
Hanna, R. 2023. "How and Why ChatGPT Failed The Turing Test". Unpublished manuscript, accessed July 13, 2023, https://www.academia.edu/94870578/How_and_Why_ChatGPT_Failed_The_Turing_Test_January_2023_version_.
Hartenberger, L. 2023. "What AI Teaches Us About Good Writing". *Noēma*, July 25, https://www.noemamag.com/what-ai-teaches-us-about-good-writing/
Hertzberg, L. 2020. "Can Robots Learn to Talk?" In S. Wuppuluri and N. da Costa, eds., *Wittgensteinian (adj.)*, Cham: Springer, 409–422.
James, W. 1985. *The Varieties of Religious Experience.* New York: Penguin Books.
Jin H., and Wang S. 2025. "Why Not Replace People With Machines? Exploring Artificial Intelligence's (In)Capacity for Education as Subjectification". *Journal of Philosophy of Education*, to appear, https://doi.org/10.1093/jopedu/qhaf087
King, M. R., and ChatGPT. 2023. "A Conversation on Artificial Intelligence, Chatbots, and Plagiarism in Higher Education". *Cellular and Molecular Bioengineering* 16, 1–2.
Lagerspetz, O. 2015. *Trust, Ethics and Human Reason.* London and New York: Bloomsbury Academic.
Li, Z. 2023. "The Dark Side of ChatGPT: Legal and Ethical Challenges from Stochastic Parrots and Hallucination". *Arxiv.org*, https://arxiv.org/abs/2304.14347.
McGee, R. 2023a "Political Philosophy and ChatGPT". Unpublished working paper, accessed July 13, 2023, https://www.researchgate.net/publication/369526036_Political_Philosophy_and_ChatGPT.
McGee, R. 2023b. "What Are the Top 20 Questions in Political Philosophy? A ChatGPT Reply". Unpublished working paper, accessed July 13, 2023,

https://www.researchgate.net/publication/369972525_What_Are_the_Top_20_Questions_in_Political_Philosophy_A_ChatGPT_Reply.

Munz, V. 2024. "Discussions of a Private Language: Wittgenstein and Rhees". *JoLMA* 5, 233–296.

Parviainen, J., and M. Coeckelbergh. 2021. "The Political Choreography of the Sophia Robot: Beyond Robot Rights and Citizenship to Political Performances for the Social Robotics Market". *AI & Society* 36, 715–724.

Pelland, J.-C., Trächtler, J., and H. Love. 2024. "Practice Makes Human: Why We Can't Understand Black-Box Artificial Intelligence". In B. Ball, A. C. Helliwell, and A. Rossi, eds., *Wittgenstein and Artificial Intelligence, Volume II: Values and Governance*, London: Anthem Press, 183–202.

Peters, M. A., L. Jackson, M. Papastephanou, P. Jandrić, G. Lazaroiu, C. W. Evers, B. Cope, M. Kalantzis, D. Araya, M. Tesar, C. Mika, L. Chen, C. Wang, S. Sturm, S. Rider, and S. Fuller. 2024. "AI and the Future of Humanity: ChatGPT-4, Philosophy and Education – Critical Responses". *Educational Philosophy and Theory* 56 (9): 828–62.

Rhees, R. 1954. "Can There Be a Private Language?" *Proceedings of the Aristotelian Society*, Supplementary Volume 28, 77–94.

Rhees, R. 1969. "Art and Philosophy". In: *Without Answers*, London: Routledge and Kegan Paul, 133–154.

Rhees, R. 2006. *Wittgenstein and the Possibility of Discourse*. Oxford: Blackwell.

Ross, C. 2013. "Spectacle, Ecstasy, and Music Festivals: The Spiritual Longing of the Lost Generation". *Liturgy* 28, 42–50.

Samarin, W. 1972. *Tongues of Men and Angels*. New York: Macmillan Press.

Standish, P. 2026. "Rather More on AI From the Point of View of Ordinary Language Philosophy". *Philosophical Investigations* 49, 3–24.

Sterling, B. 2024. "Preliminary Notes on the Delvish Dialect". *Medium*, July 7, accessed July 15, 2024, https://bruces.medium.com/preliminary-notes-on-the-delvish-dialect-by-bruce-sterling-ce68a476247b

Sussman, H. 1995. "Maxima Moralia: Millennial Fragments on the Public and Private Dimensions of Language". *MLN* 110, 856–887.

Tlili, A., B. Shehata, M. Agyemang Adarkwah, A. Bozkurt, D. T. Hickey, R. Huang, and B. Agyemang. 2023. "What if the Devil Is My Guardian Angel: ChatGPT as a Case Study of Using Chatbots in Education". *Smart Learning Environments* 10, article 15, https://doi.org/10.1186/s40561-023-00237-x.

Vallor, S. 2024. *The AI Mirror*. Oxford: Oxford University Press.

Williams, S. 2025. "Understanding Grieving for a Chatbot Using Two Concepts From Wittgenstein". *Phenomenology and the Cognitive Sciences*, to appear. https://doi.org/10.1007/s11097-025-10100-8.

Wolcott, R. H. 1970. "Schizophrenese: A Private Language". *Journal of Health and Social Behavior* 11, 126–134.

Zazzo, R. 2009. *Les jumeaux, le couple et la personne*. Paris: PUF.

Biographical note

Ondřej Beran is an Associate Professor of Philosophy and one of Heads of Research at the Centre for Ethics (Department of Philosophy and Religious Studies), University of Pardubice. He writes about various topics, often from the Wittgensteinian perspective. His recent publications include the book *Ecological Emotions* (Palgrave Macmillan 2025).