

Human Generated Text

Bernardo Subercaseaux (bsuberca@cs.cmu.edu)

April 5, 2023

During these days of computer-generated text flooding social media, I propose this small physical space as devoted to reflecting on the power of literary text written by humans. The following paragraphs aim to identify that special *something* that is present in human poetry and Large Language Models don't get, or can't get. But even if you, the reader, don't agree with them, please feel welcome to enjoy some human poetry regardless.

The distributional hypothesis in linguistics states roughly that the *meaning* of a word, or even a sentence, is nothing more than the *distribution of contexts* in which it can appear. Firth (1957) succinctly wrote: "*A word is characterized by the company it keeps*". Even if the distributional hypothesis holds water, there's a case to be made for how semantics (that which relates to meaning) can only be captured partially by language models in their current form.

Large Language Models reduce the idea of context to a purely syntactic one; in LLM-land, the context of a particular occurrence of a word is merely the sentence, paragraph, or book it was taken out of. In contrast, we humans know that in its broadest interpretation, *context* includes as well things like the mood of the speaker, their relationship with the audience, their motives to be speaking or writing, etc; all of which can sometimes leak into the syntax, but don't necessarily do. For a floral example, let us think together about what a "*red rose*" is. The distributional hypothesis says that **red** roses are whatever people buy with 15% off on February 14th, or 50% off on the 15th. It says that "red" is a color as it goes where colors go, next to "houses", and "dresses" and "carpets". It says that "rose" is a flower as it goes where flowers go, next to other flowers, next to "spring", and "bees" and "ladybugs". It says that a "**red** rose" is romantic because it goes next to other romantic things, like couples and sunsets and funerals. This whole thing gets the arrow of causality wrong: **red** roses are not romantic because they're gifted on Valentine's Day; they're gifted on Valentine's Day because they are romantic.

The surrealist Spanish painter Salvador Dalí once wrote¹ something like “*The first man to compare the cheeks of a young woman to a red rose was obviously a poet; the first to repeat it was possibly an idiot.*” This quote hints at something deep: our use of language follows from the world, and not the other way around. It is certainly possible, as we have observed with recent models, to create models of the world by its textual shadows. But let’s not fool ourselves, the red rose had an effect on the human eyes and on the human brain before we called it “red rose” or “rosa roja” or “rose nyekundu” or even שושן אדום.

If red is the color of the rose, the gala dress, the Hollywood carpet, and the Palestinian triangle, then we must wonder *what the heck is that element of the real world that makes the human mind stop more often to look and smell the red rose over the pink one? Why have multiple distinct human cultures, isolated from one another, all chosen to dye threads with that vibrant red pigment, derived from insects, plants, and minerals, even dating back thousands of years?* There must be something to the color red that makes our monkey brains go “Boom, that’s red” – that makes us want to dress in red like the most memetic flowers, and drives our attention to the regal and durable coats of the summer ladybugs. My point is, red is part of the human experience, and the word red that populates the internet is only a shadow of that experience. LLMs, who live only in a world of syntax, do not capture the extension of pragmatics (that which relates to the use and context of language), a fundamental aspect of human language. Granted, semantics and pragmatics are abstractions that leak into syntax, and thus LLMs can depict shadows of them, sometimes even beautifully so. But there is a real world outside the syntactic 45TB cave, and once we surrender into being blinded by its light, it’s hard to be impressed again by shadows and mirrors, as clean as they can get to be.

LLMs capture the distribution of usage in a frequentist manner. They capture how many times the flower that is described as red is a rose and how many roses are described as red. But they don’t capture the real underlying distribution that generates those usages; they don’t because they can’t, because it’s not hiding in the dozens of terabytes the LLM monsters eat for breakfast, it’s out there, in the actual red rose, in the actual bouncing of the light beam that goes from the cheeks of a stranger to the cones of our eyes and allows for a suitable comparison with red roses. What I mean to say is that the frequentist approach fails to capture certain important things. It fails to capture how natural it is for Tom to start calling his wife Anna by the name of “Pansyann”, because when she wakes up first on certain Sunday mornings and offers him a Cappuccino, implying to use the red Italian machine her friend Katie gifted them as a wedding present, she clearly has the vibe and the purple color of a Pansy. She’s also his Ann, and she feels seen and called in that word, Pansyann, that has never been uttered by any other human before. If you knew the real distribution where words come from, you’d know that Tom used the only word he could have used, the one that dominated the loose softmax in his brain, completely agnostic and uninterested in its prior frequency of 0. An idea often attributed to Richard Feynman is

¹In the preface of the book *Dialogues with Marcel Duchamp* (Cabanne, 1986).

that in order to simulate the motions of a falling leaf to an arbitrary degree of accuracy, you'd have to simulate the entire universe — I believe something similar happens with language.

As I was saying, my point is that we should take some time to think about pragmatics; the contribution of context to *meaning*, meaning how the mental state of the speakers, the fact that they ran out of cheerios for breakfast that morning, is part of the *meaning* of their words and sentences. My point is that we should take some time to think about pragmatics because pragmatics is the part of linguistics that captures the most fundamental aspect of poetry: that rhymes and alliterations, and the ABBA-ABBA-CDC-CDC of the sonnets of Petrarch, are just a bunch of excuses. Surely beautiful at times, but a bunch of excuses nonetheless. A bunch of ornaments and bells and whistles. In my opinion, the really really cool about poetry is that under its company we get to feel a little bit less lonely about the human condition. We have developed poetical tools that tell us that despite how alien our technology would be to Petrarch, our clumsy romantic failures and myopic quests for political power would be the opposite of alien to the Italian master. We know, for example, from the writings of Homer, that over twenty-eight hundred years ago, humans also felt tempted to abandon everything they were working for and retire to hedonic islands of carnal pleasures and nootropics and sedative drugs. We know, by the love letters of Abelard and Heloise, the French monastic lovers of the 12th century, that the desire for time to stop while your lover and you can take as many hours as you want to simply walk around, to walk around (the mountains around Puy de Sancy in their case, and the corners of Schenley park in our case) is not a fruit of modernity. Throughout the ages, we humans have felt compelled to talk about our childhoods and our fears, and our dreams about the future, and our internal monologues, and the sad soliloquy we routinely fall for on certain Sunday nights when the atmosphere feels 100 kilo-pascals heavier than usual. Poetry has served, even bypassing the hurdles of continents and centuries and languages, to make us understand that our human feelings are not that unique and that we are not that lonely after all. The space of feelings we routinely inhabit has been populated before, by countless others, and they have left us maps and candles to accompany our stay.

My hope with this project is therefore to remind the reader, whoever that gets to be, that in the lands of text generated by humans, we can find bits of the human soul and light beams of the human experience we can relate to. My hope is to showcase perhaps that part of what makes a love poem *poetic* is the embarrassing image of the poet, who sits down at night and writes the poem alone amongst their candles. Sometimes the poet fails to have the bravery required to show the poem to that one person who inspired it. But in that poem which nobody ever reads, who dies forgotten in a cabinet, there is nonetheless a snapshot of the human experience that has crystallized. If by any miraculous turn of events we, the rest, get to read that poem, we must understand something: the snapshot is not the point; the point is the thing the snapshot tries to capture and fails. Korzybski (1933) said it tersely: "*The map is not the territory.*" What I mean to say, perhaps, is that even though you might generate photorealistic images of couples in white linen clothes

that run through that blurry line where the ocean and the sand meet at night, you can't generate the couple that runs through the line, and their story and their feelings; the value of the picture of the couple running through the beach at night is at least partially due to its relationship with the outside world. It depicts something with value outside of the picture: an actual couple with an incommensurable complex life story that exists before and after the picture, and to which the picture is only a shallow pointer, a shadow.

Text generated by humans is a space for us to enjoy looking at syntax and pointers and shadows, with the promise, and because of the promise, that a real human experience is on the other end of the pointer. By following the pointers, we might arrive at that scary dimension of the human experience that others have been to before, and that way, accompanied by our ancestors, we might feel a bit better, a bit less alone.

I'll conclude by sharing a bit of poetry that ties some of these ideas together. It's an excerpt from an unexpected source of poetry: Richard Feynman.²

"A poet I think it is who once said the whole universe is in a glass of wine. I don't think we'll ever know in what sense he meant that for the poets don't write to be understood. But it is true that if you look at a glass of wine closely enough, you'll see the entire universe.

There are the things of physics: the twisting liquid, the reflections in the glass, and our imagination adds the atoms. It evaporates, depending on the wind and weather. The glass is a distillation of the earth's rocks and in its composition, as we've seen, the secret of the universe's age and the evolution of the stars. What strange array of chemicals are in a wine? How did they come to be? There are the ferments, the enzymes, the substrates and the products, and there in wine was found great generalization: all life is fermentation. Nor can you discover the chemistry of wine without discovering, as did Pasteur, the cause of much disease. How vivid is the claret, pressing its existence into the consciousness that watches it?

And if our small minds for some convenience divides this glass of wine, this universe, into parts: to physics, biology, geology, astronomy, psychology and all, remember that nature doesn't know it. So we should put it all back together and not forget at last what it's for. Let it give us one final pleasure more: drink it up and forget about it all."

²From the compilation of lectures at Caltech titled *The Relation of Physics to Other Sciences* (1964).