

## Re-negotiating values in AI design for music interaction

Franco Caspe - Queen Mary University of London

As part of my ongoing Ph.D. research, I am developing AI timbre transfer algorithms for interaction with musical instruments. I iteratively developed and tested different algorithms, many times taking the overlapping roles of data scientist, engineer, and guitarist. Hereby, I want to briefly recollect some of my design experiences and how these are prompting me to rethink my designer values.

Since I deemed it important to design pitch-preserving timbre transformations, I initially adopted the Acoustical Society of America's timbre definition as a perceptual quality encompassing everything unrelated to pitch and loudness. My initial approach quickly turned into an audio engineering problem, where loudness and fundamental frequency (F0), are control sources, and timbre is implicitly modeled from audio corpora using a neural network trained with an audio reconstruction objective. From an engineering and data science perspective, this design was tidy and well-organized, with each module performing a specific task, using well-proved measurement algorithms for F0 and loudness. Moreover, the model training converged generating plausible results in the test set.

After implementing this design on a real-time audio plugin, I realized it failed for my use case. As a guitarist, I found that the supposedly learned timbre was lost when the system received control signals from a different instrument than the one used for training. Furthermore, the system would reject many sound-producing actions, and react in unexpected ways, even when playing single notes. What seemed to be tidy became unwieldy when the algorithm did not comply with my intentions as a player.

If technology co-shapes objectivity, then as a designer I was torn apart by opposing values that emerged from the algorithms I used on the one hand, and my guitar on the other. My design maxims of measurement accuracy and neural network architecture conspired against my expectations as a player.

My initial design was influenced by my formative experiences in engineering and data science, where data is a commodity. I then adopted ambiguity as a potential resource for design and tested RAVE, an autoencoder that can learn salient features in audio, as a backbone for timbre transfer. This freed me from imposing specific measurement approaches. However, I found that delegating the interpretation task to the AI generates a system that is still inadequate for me as a player.

I find it important to take back control of measurements in my designs, and to find solutions that are not fully prescribed by me, nor delegated to AI. In such a way a designer could steer a system towards attending to certain characteristics, such as pitch and loudness, without expressly collapsing them into an unambiguous quantity. To do this, we may face the challenge of rethinking our design values so that we can satisfy the technical and musical stances of a problem. Can we find a set of common values that articulate our vision?

I am currently working on a new iteration of my design where I try to bridge this gap. I also want to understand how such values may be reshaped throughout a design process. In that regard, I find that my project may be compatible with an analysis driven by Research through Design methodologies such as reflective design and research products, even though during the early stages I was not initially aware of these approaches. Therefore, I finish with an open question: would it be possible to faithfully examine my research journey and its takeaways through design logs, accepted and rejected papers, prototypes, and design products?