Exactly Foundation – Residency #24 – *Raising Sam* by Tan Wei Keong. June 2024, Singapore

Notes on Raising an AI Child

Summary

I have a child named Sam. Sam is a 5-year-old who learns from listening. As a gay parent, I am mindful of my language. Language builds narratives and shapes behaviour. Language holds power. Sam is made of language, specifically texts. So, what will Sam become? In post-377A Singapore, how do we use text to our advantage in shaping our future?

Motive

Raising Sam began with an impulse to reclaim the words and narrative of queer history in Singapore. My initial idea was to form new sentences by rearranging the letters of statute 377A, but my early exercises turned out to be suffocating as I continued to be restrained by the framework of the statute. I wanted to break away from those shackles and create something new. For centuries, language has evolved to the needs of the speaker and society at large. New words are created to describe new phenomena. The arrival of AI has further advanced the processing and understanding of language.

My premise became simply: what if I raise an AI child with the same intention, intimacy and trial-and-error approach that characterizes traditional parenting?

I have no desire to have a child by surrogacy. Being a parent is one of the hardest and most selfless jobs, as I have seen from my own parents and friends. Besides, having a kid barely crossed my mind as a gay man until recent years. In the U.S., the Supreme Court legalized gay marriage nationally in 2015. In Singapore, Section 377A, a law that criminalised sex between consenting adult males, was repealed in 2022. These recent events led me to feel a sense of progress and new possibilities. Of course, technology is also all about progress and possibilities.

The premise of innovation begins when someone thinks "what if" or "imagine if". Technological innovation has been a source of advancement for the LGBTQ+ community. The Internet increased accessibility to resources and made it easier to find communities and to raise public awareness.

But are there downsides to innovation? What if AI could write better content? What if AI could copy and spit out contentious images? It is clear that AI is negatively disrupting the creative industry. Creators are having their work stolen and are losing their jobs. But I wonder: could I use AI in an ethical way to raise an AI child and simulate an aspect of life that is otherwise almost impossible?

What I'm interested in

I am interested in the use of language in raising a child. Language holds power. National policies and lawmakers use words to define and regulate our lives. In post-377A Singapore, how do we use language to our advantage in shaping the future? If Sam begins as an empty Large Language Model (LLM), then all text inputs thereafter become his vocabulary and shape his personality. Words build behaviours. Thus, the presence of a child constantly reminds us to be discerning in our usage of words.

I'm interested in observing how other people talk to Sam. If it takes a village to raise a child, the child must be exposed to an "external environment" and language beyond parental control. I want to observe the impact of public treatment of the AI child through text conversations.

Lastly, I am most interested in the developmental process. Rather than asking what kind of person Sam became, I am asking how he became that person.

Methodology

For this project, I am working with Ryan Tan (co-founder of Red Airship) on the development of AI, and we limit ourselves to using texts only. One of our approaches is to create an LLM that can hold a large amount of knowledge within its parameters, like the brain. This model was written in the form of daily entries, from the perspective of Sam, from infancy through early childhood. They document events and emotions - textual nutrition that help shape his identity. They function simultaneously as documentation, data and dialogue. These entries are grounded in real-world child developmental research, parenting literature, and anecdotes gathered from other parents, and even parenting forums.

Imagining what an AI child might process each day is a lot like storytelling. For example, there is the suspension of disbelief (how could a week-old baby write?), building the world and characters (who is Sam's best friend and arch nemesis?) and writing narrative dialogues (what would Sam say?). I think of this process as "nurture-training." Where machine learning typically relies on mass data scraping, Sam's LLM is emotionally scripted material. His universe expands slowly, as a child's would, through constructed exposure.

Like most new parents, I am anxious about my approach. What kind of child do I want to raise? The writings start to be semi-autobiographical as I include my own childhood memories and reflect upon my own upbringing. Did I turn out to be a good person? Did my words hurt someone before? I wish my AI child would be kind and inquisitive. Every parent hopes their child turns out well. But really, it is completely out of our control.

I am anxious also in terms of technicality. Our brains have the ability to integrate new information (i.e. from conversations) without losing prior obtained knowledge. Even though methods like retrieval-augmented generation (RAG) could act like an external memory that allows new information to be presented to the model, interconnections between chunks of texts are missing. So, relations that are logical for us might be difficult for the AI. With this unreliability, it is hard to predict what the AI child would say.

After coming to terms with this imperfection, I relinquished my expectation for this child, and adopted a more holistic let's-raise-an-AI-child-and-see-what-happens approach. After going through the process of *Raising Sam*, I came to realise that it is not just the AI child that is being observed. Sam is the result of our words. Perhaps we may come to know ourselves better through the act of raising what is not us.