Nairobi River
Kirinyaga Rd
Kilome Rd
Jevanjee Gardens
Koja
Moi Avenue
Tom Mboya St
Latema Rd
River Rd
Voi Rd
Kirinyaga Rd
Accra Rd
River Rd
McMillan Memorial Library
Jamia Mosque
Banda Cl
Kimathi Street
Banda Street
Kenyatta Avenue
Moi Avenue
Luthuli Av
Kenya National Archives
Moi Avenue
Tom Mboya St
City Hall Way
Holy Family Basislica

Start / End

Walking Route
(ca. 2.3 km,
ca. 29 – 40 min)

You will find sounding encounters along the way. Some might be more obvious than others. Some are planned some are not.
Enjoy your time!

ear map

AI-Soundscapes

by Brian Muhia

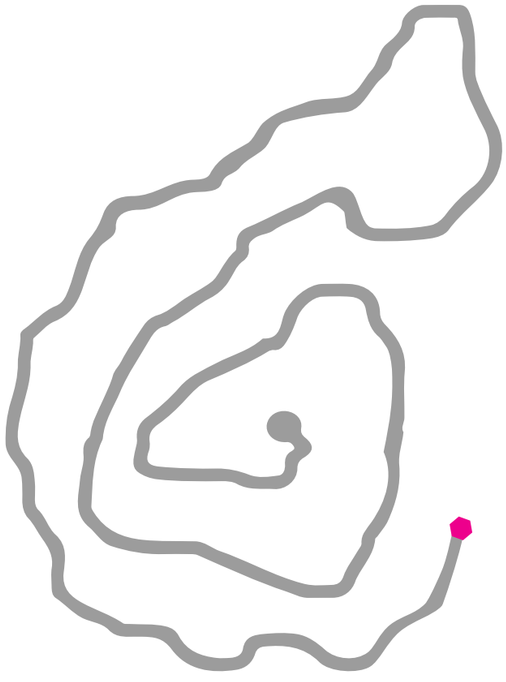

This is an experiment: Sound of Nairobi is a archive of field recordings that exists since 2019. By now we have over more than 100 recordings in the archive from various parts of the city, made by different people over the past years. Now I ask you to consider treating the entire archive as an environment of its own. A closed space where you may take a virtual walk through different auditory experiences that the city generates.

We trained a Real-time Audio Variational auto-Encoder (RAVE) on some samples of the Sound of Nairobi archive of recordings. This is a method for learning a model (in the sense of Judea Pearl) of our environment that we can play back for ourselves and ponder the state of our world from a different perspective. This computer generated environment of learned representations by a system is glitchy, sometimes producing random silence, or playing a bird's song only to rumble with the momentum of a truck in no particular order. But isnt world sometimes exactly like that?

You are on a walk through the city. Some sounds are familiar, you can decode them you know where they come from and what produces them. And then at the next moment you hear something you never heard before, sometimes if you are willing you can find out what that was, sometimes it remains unknown to you. Are these glitches? And then out of some reason the volume of the city drops and theres is something like silence ..

The source of sounds in the archive and what we trained the model on are binaural recordings made by collaborators and volunteers. The binaural recording technique is a way to record sound around you similar to the way you hear sound. While recording you wear two small microphones in your ears, like headphones. The microphones cathch the sound at your ear, exactly where you hear sound. It is a pretty realistic way of recording the sound of your environment.

These recordings are now the basis for new AI generated soundscapes. Although this recording technique seems very objective and the recordings thus one might claim are a record of reality you cannot strip the subjectiveness of the recordist from them. It is a person wearing those microphones with all their specific bodily traits, their height determins where the mics are, their gage determins how fast the recording is made in one space, if they turn their head because something is grabbing their attention you can detect this in the recording. This is he beauty about these recordings they are inherently subjective and a rendering of one persons reality.

Thanks to research from LAION we will soon be able to type text and hear a sound that matches what we typed. This is part of a cultural shift for AI development which lets communities and individuals choose what and how to share with the world, so we are open to discussions about the most ethical approaches towards releasing tools and systems based on representation learning. In 2019 the state of research on audio generative models was vast with a long history in text to speech systems, and it has only gotten bigger since. Our choice of techniques is based on systems with available source code, and it is thus reproducible and possible to collaborate with other teams on open source software.

by Brian Muhia