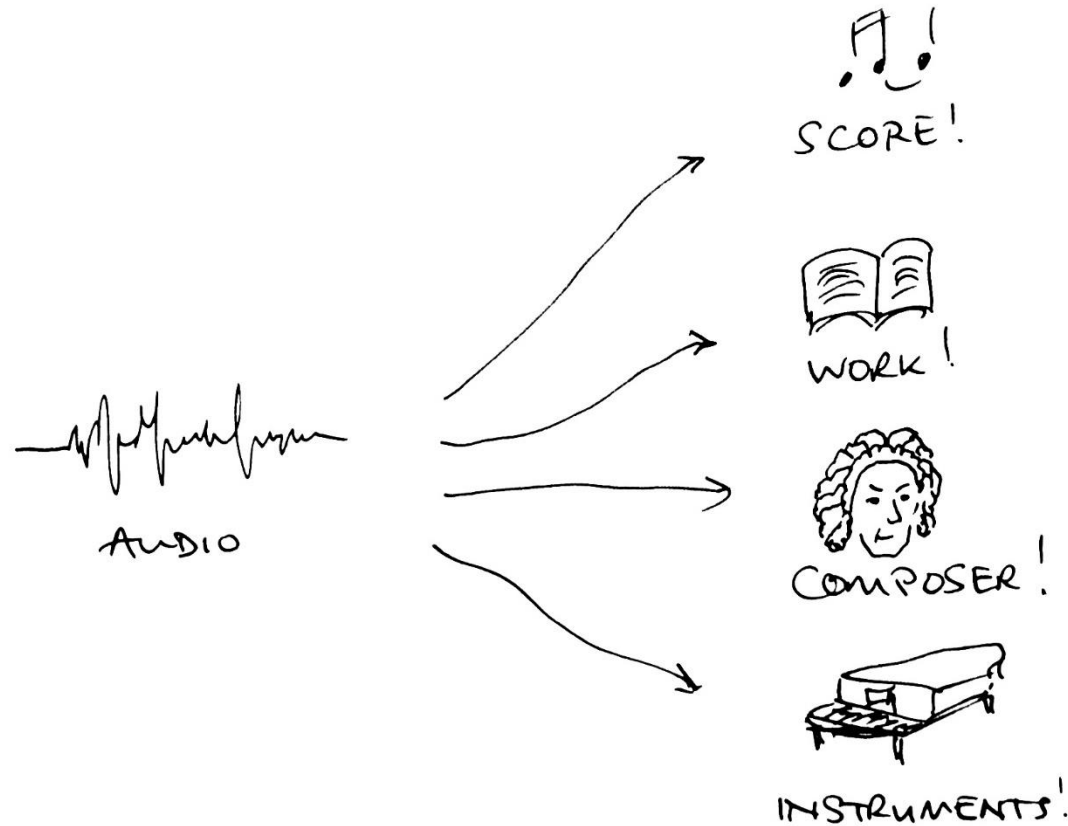


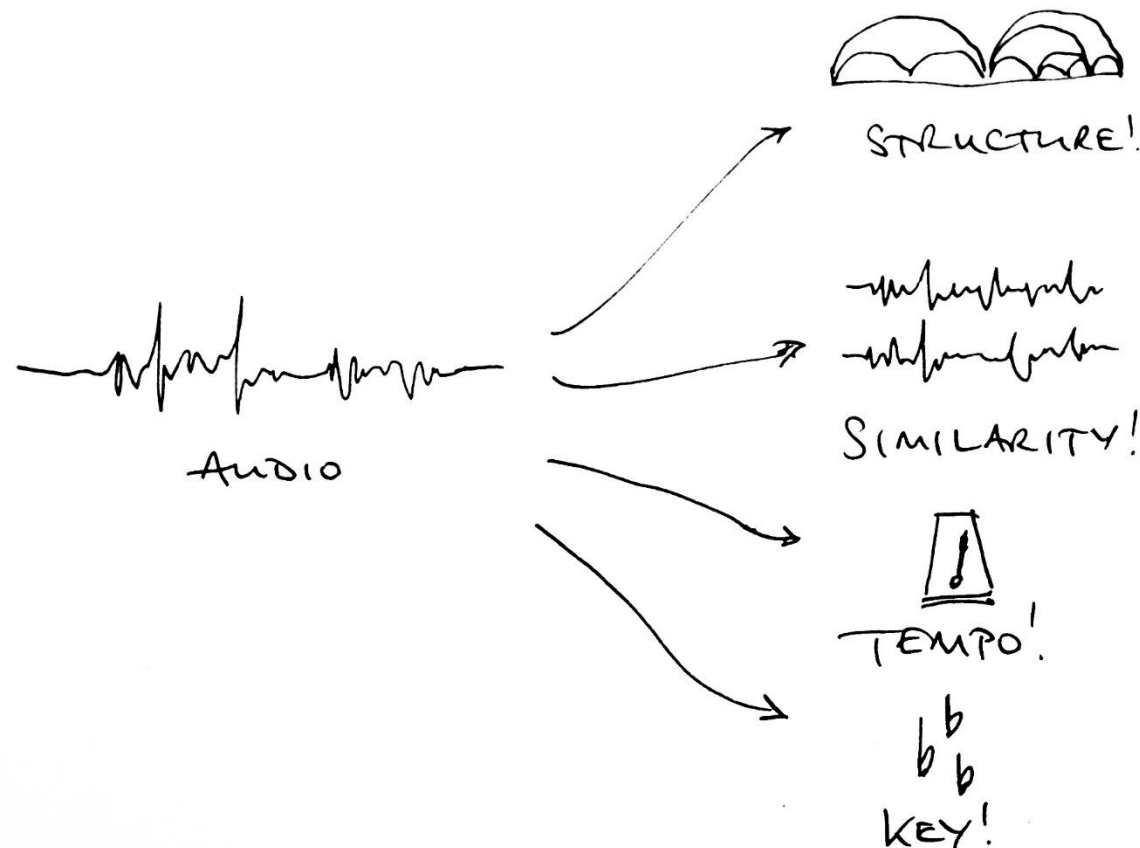
What do we mean by “audio feature”?

Ideal: TRUE MEANING extracted from the audio signal




What do we mean by “audio feature”?

Ideal: TRUE MEANING extracted from the audio signal



What do we mean by “audio feature”?

Reality: something we can squint at & interpret a bit



“Low-level” and “high-level” features

Low-level: “mechanically recovered” from the audio

- e.g. amplitude, timbral descriptors, spectral features

High-level: usually obtained from low-level features + lots of context (template matching, machine-learning, domain knowledge)

- e.g. key, pitch, tempo, notes, phrases, similarity

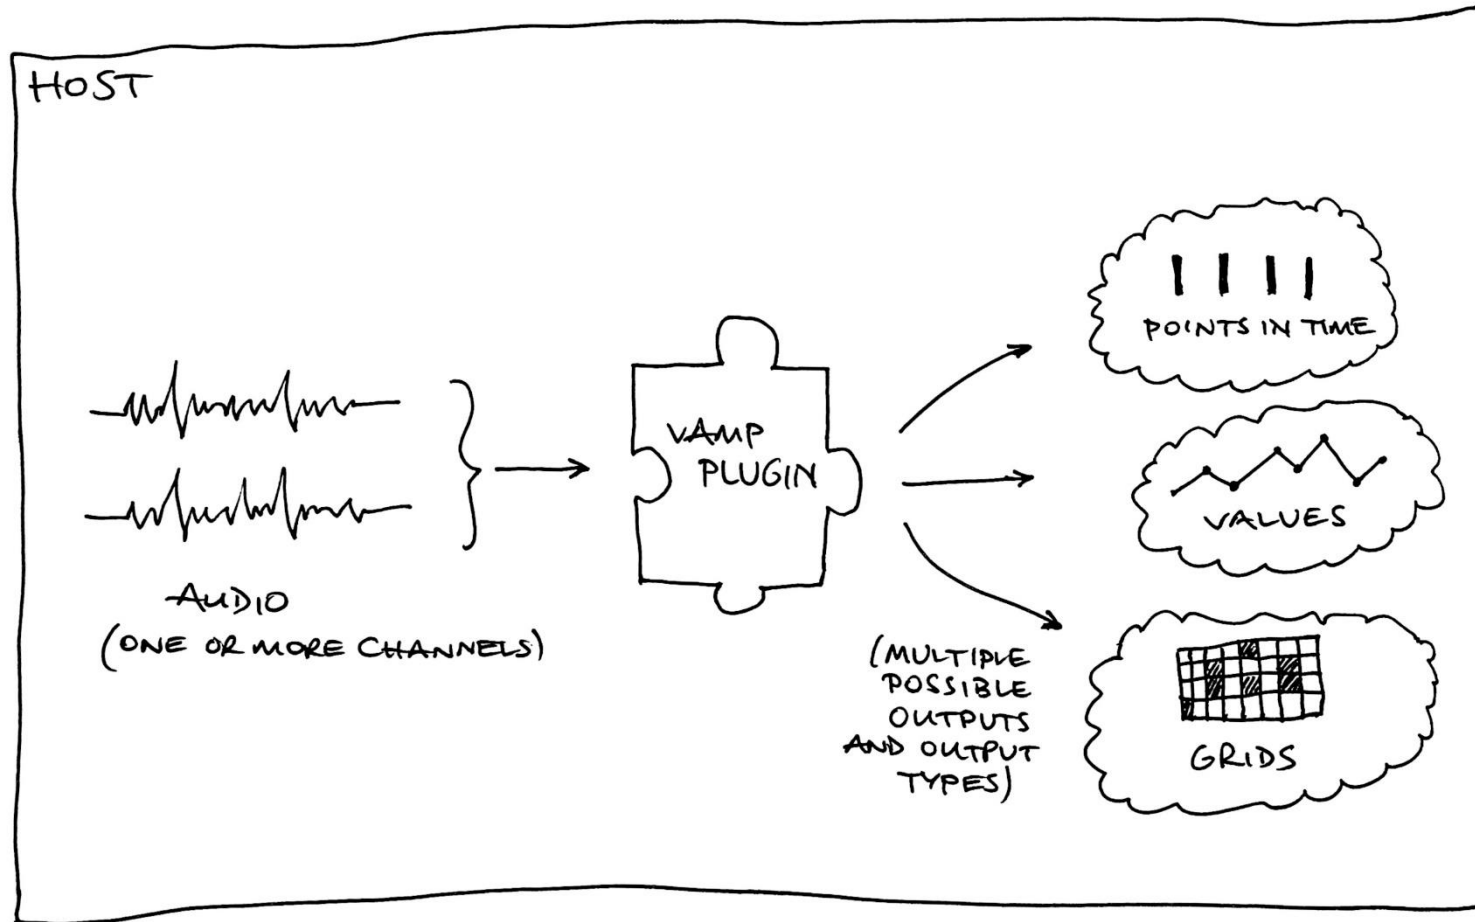
Vamp plugins

Small files you can install that add new feature extractors.

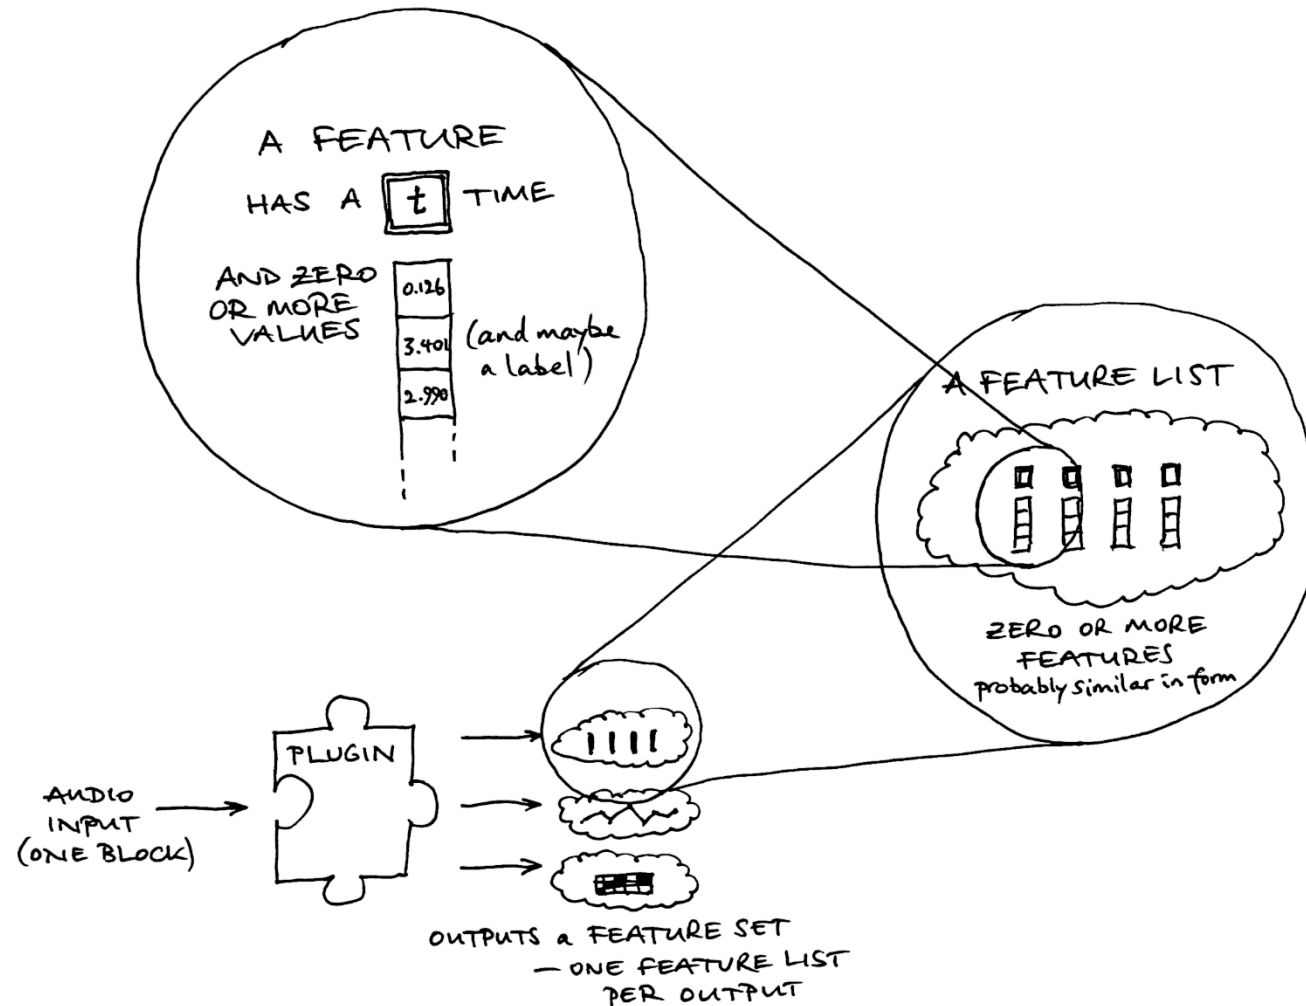
Once installed, can be used with several different “hosts”:

- Sonic Visualiser
- Audacity audio editor (simple feature extractors only)
- Sonic Annotator – batch audio feature extraction program
- Python Vamp host – use with scientific coding packages for analysis, search, plotting etc

Vamp plugins and audio features



What does a Vamp feature consist of?




Example: Chromagram

Somewhat representative of time-varying harmonic content

- Made by “wrapping around” time-frequency spectrogram into a single octave
- Various ways to do this → lots of different chromagram plugins

Good example of an *almost* intuitively meaningful feature



Chromagram

Motivation

- Reduce spectrogram in a way informed by musical structure

Limitations

- Time/frequency resolution tradeoff
- Misleading outcome of harmonic folding (different approaches to this)
- Intrinsic difficulties, e.g. with temperament

Applications

- Chord and key estimation
- “Harmonic feature” for search, retrieval & similarity tasks