

Intellistent Audio Archiving

Research hypothesis:

An artificial intelligent (re-)combination system in an audio environment will augment the listeners ability to get specific information out of an (audio) archive and will afford (radio) makers new possibilities and insights on how to make a "program".

Goals:

The direct goal is to make an initial step by delivering a very basic application and gathering a community around it to support and test that basic application and create the conditions to propel it forward.

The more distinct goals are in developing an artificial intelligent conversation partner. In that sense the intended "dreamed" result is to deliver a system to the end users which allows them to listen to parts of an audio archive and browse through that archive in an intuitive way. I want users to have the same feeling while moving from thread to thread as when they are in a conversation with somebody and the conversation "flows" nicely: they have the feeling that this is dealing about exactly what was on their mind for the last day/week/month and now finally somebody is talking about it or discussing the topic at exactly the same level as they are. The feeling of singularity, of convergence, that everything fits perfectly, the euphoria of standing on the front top of this (mind)wave they are surfing.

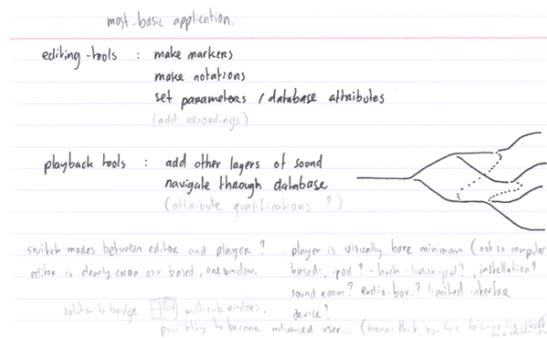
Explanation of the term "Intellistent":

Intellistent is a derivative from intelligent, which is Latin and comes from inter (in between) & lego, legare (to collect, to gather, to read, to read out loud, to speak justice and to overhear). So intel+ligent actually means to read in between the lines, or to select the best from a collection. Note that the dutch word "lezen" (read) also means to gather (this is what you do with the ear of grain to select it, you "read" it). And "uitgelezen" (read out) also means selected. Hence the meaning of the word intelligent: the ability to read from in between the lines, to get more information out of a text than what actually is written on those lines. Analogous to intelligent "reading in between", intellistent is a combination of Latin and english and stands for listening in between. (In fact "intelligent" itself could already be translated as listening in between, but it would be a bit hard to get this translation out of the word, because of the commonly addressed meaning to it. Furthermore one could also think of "intelaudent", since this would be more native Latin. However I feel "intelaudent" is less recognizable as related to "intelligent" and I like to make a distinction between passively hearing (audio, audiare) and actively listening.) Intellistent is besides related to a "smarter" understanding of the spoken text, also more physically related to the project in which I like to weave multiple threads of audio together. And a combination of these could only be understood by listening in between the different audio layers, to get their connection, to understand the commonness of the threads.

Description of the most basic outcome of the project:

An application, running on Apple's OS X, in which a user can record and listen to audio and which allows the user to make annotations at the same time. This application uses microphone, keyboard and jogwheel/jogdial input. The application has two distinct modes of operation: recording and playing back. It would not be desirable to add recording functionality to a playback mode or vice versa. In the recording mode the standard keyboard would better be replaced with a keychord for more flexible application: you don't need a laptop in front of you. The application allows for linear annotation: the making of notes and remarks along the timeline, while the audio is playing (or while the recording is being made). When the user restarts typing again (after a pause of a few seconds) the application makes a new marker on the timeline and the typed text is inserted from this point.

Together with the ability to save and open files and scrub (rewind & forward) through them, this concludes the functionality of the most simple application as the first result of the project.



sketch. The threads illustrate a possibility tree of navigation through the audio archive in a non-linear way.

Right after this initial step a few enhancements could be immediately thought of: an application that allows users to choose from multiple files and with which one can attribute certain characteristics and values to these files: who, where, when the recording was made, how long the audio is, a title, a summary, and maybe some keywords or some other valuable signifiers.

Another primary addition could be multiuser functionality. Allow multiple users to access and annotate the same audio files. This would have to deal with a system of permissions for users and a way to check files in and out. Some wiki-like or CVS-like functionality. It is also interesting in this sense to have a look at how SubEthaEdit(1) is adding multiuser functionality to something as basic as a text editor.

This basic application and the ability to demo and test it is critical for the projected continuation of the work. Therefore the launching of version 1.0 is planned early in the overall project timeline. After this launch user testing and evaluation can begin and I can (hopefully together with a group of early users / believers) think of how exactly the system and the application should be continued.

Future projections:

As said an audio annotation tool is still far away from an intelligent conversation partner. However I have given the road in between some thought and have a couple of techniques in mind that could lift this application from a very simple linear playback tool to a much more complex artificial intelligent navigation and recombination environment.

I would like a system to be researchable by agents (bots), that try to find interesting leads and combinations independent of (but also related to) the users browsing behavior. I could see a future in which the material works as DNA strings that can be automatically mutated and recombined. I think the concept of memetics fits closely with this process. And when there is a growing amount of users a collaborative filtering system that gives recommendation becomes interesting. Maybe then one could browse associational through the system, starting with one keyword. But I am also thinking of describing the audio along a multitude of axis and thus creating a multidimensional (virtual) space in which the audio traverses following the specific multidimensional values. Each part of the audio would have a particular direction in this multidimensional space. And because it could also still be cut up and recombined in interesting ways, the multidimensional space becomes filled with little threads that all have a particular direction and from which new threads and pathways could be formed. This way a rhizomatic structure would emerge out of the audio archive -or you could say: grows into it- almost topological, like crystals are formed in a salty evaporating solution. Such a threaded structure with internal directions could even enable preemptive browsing, predicting the route that the user is most likely to follow based on his previous route and the internal topological directions of the system.

However note that it is not to be said that those artificial intelligent systems will actually be build in the application during the running of this project. One should rather see these as projected directions in which the project could possibly continue. And maybe some of those techniques could be tested in a very crude way to observe how they behave and operate in an audio archive environment. So as the very first stage of the "intelligent conversation partner".

Modi operandi:

While there is a difference between the use of the application while recording and while playing back, I see actually a quite bigger difference between using the most basic application described above to edit and annotate the audio and an audio browser which would work with the artificial intelligent systems in which directions I like to see the project heading from the previous paragraph.

While the annotator / editor is provided with a wide set of tools to set all kinds of characteristics, the end user / listener is given only a limited set of interaction possibilities. This is for several reasons: I want the listener to focus mainly on the audio rather than on the interface. I think novice users don't want to learn a system, they just want to browse through the material and listen what's in there. Hopefully a limited and clear set of interaction draws the user to use this interaction to its extremes and gets the most out of the system. I don't consider zero interface as a real solution, but a subliminal interface would work a lot better than a bunch of bells and whistles. I would dream of an interface that consists of a camera system that looks at the listeners face and analyses its subject facial gestures, but I know that would be just adding to the difficulties. As a simple user input I like the dial-plus-one-button interface of the iPod. Maybe I can make something similar, allowing the user to navigate through the archive and giving ranking. Or just use a slider and two buttons.

Also the visual feedback should be extremely limited compared to the feedback of the annotator/editor. To perform tests on those AI systems and to set their properties transparency of the underlying system is essential. The editor needs to see the guts of the system, including that of its internal agents. But the end user is -I think- not very helped with layer over layer of this kind of complex information. They should be better off with a simple and clear graphical feedback of what is happening and where they are in the archive.

The only aspect of the feedback which will not be simple is, of course, the audio. While I know this could be a whole new research subject all together, I still want to mention this: I think the system would convey the topological structure of the audio threads the best if these spacial dimensions were exaggerated in the audio itself. Surely there can be multiple voices (multiple layers) mixed together. And it can also be done in such a way that those different voices are still separately conceivable, through extreme panning between left and right channel or difference in timbre, the cocktail party effect(2). I think this effect can be exploited and made to use when it comes to navigating through a complex audio archive. This would however include a lot of study of realtime computerized filtering and digital sound processing (DSP). I would have to resort to external sources of knowledge for research in this field. And while I find this field highly interesting -note that the professional movie industry has also discovered how to exploit its power-, still I do not know whether I should already invest time in these complex subjects: the processing and perception of audio.

The users:

The system will have three distinct layers of users: A. the maker, designer and programmer of the system (me, hopefully with some help). B. the editor, radio maker, documentarist, the hyper- / super- user (a selected set of test users to begin with). C. the end user, audience, listener, novice browser (the rest). The group of hyper users (B) will first be invited by me to test out the system. These early adopters then later could themselves invite other audio people that could be interested in using and testing the application. I expect this group of users will not be so big, for not everybody is a editor -though I wish they were-. Also, I like to keep the system open and transparent. Therefore I think that end users / listeners should have the possibility to become hyper-users themselves.

For the audience user / listener (C) I picture somebody with little or no prerequisites of computer usage, so this could be literally anybody (that is not deaf nor blind -or maybe that could even be- and has the ability to operate a knob and/or button). These users will only be using the system near the end of the project, when it has been through some iterations of development. In the meantime the test users (B) will have to play their part, imagine they are still fresh and ignorant to the system, or find other "lower-level" test users in their own surroundings.

While there is a significant gap between the end-user(C) and the editor-user(B) in terms of usage, interface, access and control, the end users doesn't always have to stay end users. I think it would be actually very interesting for these end users(C) if they had the possibility of gradually becoming editor users(B) as they use the system more; that they could upgrade their user status as a reward for getting involved in the project. For I reason that at some moment these end users(C) will become "the real database" and I would better try to get this valuable data in.

The content:

As I described the basic application (version 1.0) will just be for editor/annotator users and so the content used with that application will mainly be material found or originally made by these early users. With the application becoming multiuser, the users will have to upload their audio-material to a shared server on which also the annotations will be stored.

But while I was conceiving this project and was developing it in its very early stage, I was listening to a particular dutch radio program: "de avonden"(3). And it was actually this radio program that sparked my interest in developing an intelligent audio archive. All the episodes of the avonden for the last 4 years are online in real-audio , but there is no convenient search algorithm, the episodes are not linked or referenced to each other and even if they would it would just be a reference to a two hour real audio stream. So finding the interesting bit in this two hours, is still up to the internet listener and one can imagine real audio doesn't really provide a nice interface to find this needle in the haystack. At the same time the thematics of all these episodes are closely linked and mutual reference is a common thing: e.g. in a conversation some philosopher cites a writer, with whom also has been an interview in the past; or a topic like the contemporary state of crisis is recurring as a main thread in different interviews throughout different episodes.

All this led me to the idea, that a system could be laid upon the current database of episodes, which would allow an interested user like myself to browse this database in non-standard, intuitive, associational way. So I could make my own episodes. This idea -which is much more easily thought of than actual developed- led to this project.

While this project became a lot more than just an overlaying interface to the archive of the avonden ... or actually it became a lot less to begin with i.e. a simple audio annotating tool... I am still interested in using the avonden as content and maybe some of its makers as test users. The first reason for this would be a research methodological one: in order to analyze the algorithms used in the (artificial intelligent) system I think I need a consistent set of material with small mutual differences. Then every outcome can be closely studied. The second reason would be that in the dutch audio landscape this program seems to be one of the more easy paced, friendly environments, where the makers and the listeners still take their time to listen, think and contemplate, which is I think one of the required ingredients for a good conversation. Finally I like to use the content of the avonden, because I like my conversation partner to be influenced and shaped by the same set of subjects and ideas as I choose (i.e. the subjects and ideas brought about in the

avonden).

In the initial stage of the project I could just use the captured (hijacked) audio stream from the website of the avonden. Later I might try to get direct access to the whole archive. When I would use this material in my final application which is also suited for the end user(C) it might be a good idea to limit the experience of this user to also two hours and to provide a similar intro and outro to it, like the original also has.

Scenario Matrix:

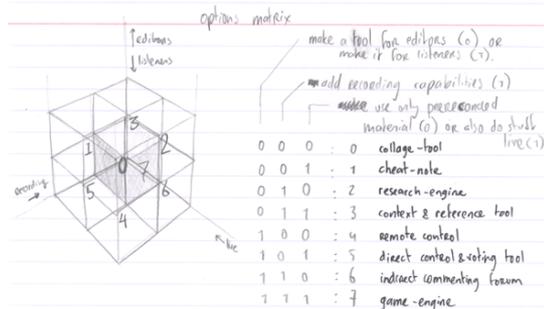


diagram of the scenario options matrix: a 3 dimensional cube with a major functionality decision on each of its axis.

Even already in its primal basic stage the system and its application could develop in a few distinct directions. I narrowed down three major decisions about the context and functionality that has to be made for determining this direction:

- Will this be a tool for editors only or do I open up the applications for listeners as well? (and thereby limiting its functionality and interface for reasons described above, see Modi operandi.)
- Do I add recording functionality to the application and the system, so allow also new material to be entered in the database through this very system, or is it only a playback tool, a functional shell around a predefined set of material?
- Is the system only to look back at prerecorded material and does it do its analysis afterwards, or can it also deal with live input and change its behavior according to this continues stream of user feedback?

These three binary choices (It is either the one or the other) make up 8 (2^3) different possible outcomes:

0: a collage tool

An automated montage system for editors / directors, which allows them to make new edits and versions of the same collection of material.

1. a cheat note

An immediate and easy access to the audio archive which would allow an interview host to use the system as a cheat-note, so he can check previous interviews for similar leads and threads.

2. a research engine

A tool for researchers, members of the editorial staff of a program with which allows they can browse the archive, annotate it and add new audio to it.

3. a context and reference machine

A tool for interviewers & the interviewed which they can use to let their conversation be steered or find reference material for posed statements. Allows them to pause the playback and react to it.

4. a remote control

Gives the end user, listener some tool to zap through the archive. Because of limited feedback it will merely bring diversion and be entertaining.

5. a direct voting & control tool

Will allows a listener to engage much more active with the "radio"-program and steer in which direction the interview or conversation should be heading.

6. an indirect commenting and feedback forum

This will be a system to add interaction to radio for the listener. The type of interaction that creates a dialog with the makers and other listeners. They can say stuff back to the system or make their own interview or documentary. It will allow listeners to give comments and steers the makers in an indirect way.

7. a game engine

When everybody can give feedback in a live and direct way, a special sort of game could be played (if it doesn't fall into total anarchistic chaos, that is). Multiple interviews and conversations take place simultaneously. The listeners can switch between different threads and can steer these different conversation in diverse or conversing directions (i.e. they can decide whether the talk is drawing towards a mainstream or instead specializes into a niche topic). The listener can even ask certain speakers to focus their thread into a different sub-topic or answer other questions on the side. Talk radio meets IRC.

The basic application that will developed as version 1.0 will be much like the research engine, 3. After that versions launching this scenario matrix should be studied again to see which technical options are possible and whether the projected result is a desirable one. The last option as example would be fun to witness and maybe also fun to participate in, but I doubt it would enrich and deepen conversations or augment the listeners ability to get valuable information out.

The Artificial Intelligent systems:

I tried to make a survey of possible artificial intelligent systems that could later be implemented in the project. I came up with four distinct techniques, but this doesn't mean that eventually there have to be made a choice which one to use, if at all one will be used. These artificial intelligent systems can also be merged again and techniques or aspects could exchange between them.

1. Genetic algorithms & memetics

A system using genetic algorithms with human appreciation as competing environment. Genetic algorithms is a computational technique in which evolutionary systems are used to select the best fitted algorithm for a particular problem. Key in this process is a limiting constraint that brings certain algorithms to extinction and lets others flourish and evolve. In the technique for using genetic algorithms I propose here, the human appreciation would be taken as this external constraining factor, the competing environment. The system will begin with random recombinations of audiopatterns. Then gradually the user learns the system how to best please him (by making sensible audio compositions, preferably). I imagine the system will learn both where to cut the data (good in and out points) and which data belongs to each other (similar threads or lines of thinking). The system could use a neural network to have this composition made. I also like the idea of using the audio pattern itself as direct input for these genetic algorithms, or even as algorithmic key itself. So for the system there will be no distinction between code and material. I believe this would be similar to how DNA is recombined and mutated inside the cell.

- + could already work for one person as user group
 - + string cutting and pasting with audio resembles gene cutting and pasting (like DNA), so it could possibly be a very accurate system.
 - includes complete audio pattern analysis. This is both difficult as advanced mathematical topic and
 - would rely on very intensive computing power.
- examples: biological human genome research, finding advanced medicines

2. Collaborative filtering (CF) recommendation

Collaborative filtering is an artificial intelligence technique in which agents work together to find similarities between groups of users and their preferences or behavior. Users are asked to log in and rate the elements of the database known to them. Based upon these (shared) ratings, the system suggests database elements to the user, that might also be rated positively. When a bunch of users all rank a specific set of elements as nice and you have ranked a (smaller) subset of these as nice as well, the system will recommend you those "missing" (still unranked) elements: "you might also like these items."

- + doesn't require special annotation or analysis of the audio data.
- + only asks users to indicate their appreciation per specific element
- requires quite a big user group to deliver preferred fine-grain results
- only known to work in areas with separable elements (songs, CD's, movies, books). No indication if it could also work with streams and patterns of material (e.g. continuous audio feed).
- "Accuracy can be a problem."

You must manage people's expectations. At first, people find recommendations cool and addictive, but after a while the novelty wears off and quality in terms of accuracy become the key focus. Does it really get better and better? The more you rate, the less certain you are about the items that you are rating, so your profile actually gets more blurred and recommendation quality can decline."⁽⁴⁾

examples: Amazon, former system of 24HCinemaService, a dutch movie recommendation site, hosted by the VPRO, the Firefly system (Patty Maes)

3. Keyword searching and organizing

When all audio material would be indexed and annotated by some editors, its relatively easy for a computer system to deliver an "associational" navigation through the archive. All audio pieces and separable paragraphs has to be annotated in a "standardized", common language. A computer algorithm operating behind the screens can then use these keywords from the annotations as similarity points and can rearrange the data so it becomes grouped by one or more of these keywords.

- + relatively easy programmable algorithm
 - + easy testable system (works already with a single user and a small subset of data)
 - no feedback of the end user is taken into account with this system
 - a huge amount of labor will be involved with annotating all audio fragments
 - the system has already been used in works of teachers/staff and former PZI graduates.
- examples: Boston Highway Movie Database (Michael Murtaugh)

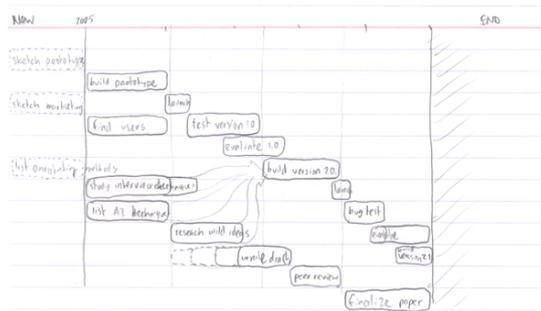
4. Multidimensional space of Key-quantities

Instead of keywords, keyquantities are used to annotate the audio feed. A strict set of key quantities is chosen. (e.g. open - closed, lively - death, organic - synthetic). A selected group of editors is asked to judge the audio feed according to one of these quantities using a slider. Based on these (multidimensional) values a multidimensional map could be made in which all audio is drawn as a line in time, traversing through the (virtual) space with the axis as key quantities chosen. Self organizing map algorithms can be used to flatten out this multidimensional map to a linear one dimensional map in which a listener can go forward and backwards, following a specific thread. The interesting part of this technique begins when the system is set up to preemptively calculate the topic direction which the user is pursuing. For the system knows where the listener has been (in this multidimensional virtual space) and it knows in which order the listener has visited these places. So thereupon the system can create a one dimensional perspective, which shows the overall direction in which the user is traversing and use this information to preemptively present user with this target ahead of time, almost like reading your mind.

- + multiple users together make a reliable average value. It can even have an accurate significance value (how important or strong is this value at this moment?)
- + can "enrich" a audio stream on a profound level. Can work with continuous feed of material. This could even be the solution to apply CF techniques on a stream. (Since again patterns could be found of certain subsets of users that have a particular shared option).
- the choice of the key quantities (in the beginning) is strongly determining the outcome (in the end). Who is going to decide which quantities will be used? Because of this it's not very flexible or open.
- to be accurate -and have the same audio analyzed by the same person on different quantities (in different dimensions)- is an incredible workload and will be difficult to ask of a user.
- works a with very complicated complex theoretical model (relating to Deleuze's Mille Plateau). Unknown how this will result in a practical application.
- the system generates a lot of values. Which ones are most relevant and which ones less? Or which ones are relevant in relation to which? The system will have so much floating values that analysis of its behavior will be increasingly difficult.

examples: datamining, perception analyzer (voting tool to analyze the strength and weaknesses of a particular debate or speech), Datamap stich & split (<http://www.stitch-and-split.org/site/viewer/preview.php>)

Time plan:



the time plan runs in two main lines: the programming, launching and testing of the application part and the theoretic research. Halfway through the project some theoretical issues could be to be merged with the application.

Note:

Launching the application (happens twice: version 1.0 and version 2.0) will also have to include some marketing and user manuals to be made public on a website. This involves maybe just a little more work than is indicated on the schedule here.

Update on the time plan (12-12-2004):

The time plan above has two -sort of- independent paths leading to a work and a thesis. At this moment it has become apparent that a third path have to be added to the time plan. That path will deal with external communications. What to tell who and how to tell that.

External communications:

In this project user testing and user involvement (participation, co-designing(5)) will be one of the foremost success factors(6). That is the reason why external communications will be a major part of my final project. The main medium of this communication will be the project website. This online resource will -besides offering downloading the beta versions of the application- be the place to get informed about project planning, immediate goals and long term objectives, recent changes and will feature a platform for discussion and most important: the place where a user community will be formed.

One of the main texts on this website that will contain this project information will be a white paper.

The topics of the white paper will be:

1. Introduction
2. Summary
3. Vision, Mission, medium and long term Objectives and Strategies
4. Present Status of the project
5. Project values, Terms of Use and Licensing
6. Interim Goals and a up-to-date Time Schedule
7. Functional description / detailed requirements definition of version 1.0
8. Profile of Target Markets and User Categorization
9. Organization Structure (and who is cooperating, and who are we still looking for?)
10. Conclusion

Besides this formal white paper, all this information will also be made available in easy readable dutch.

Technical issues:

(Besides the AI.)

When I get access to the complete archive of "de Avonden" and want to make that all available in a networked setting, it will be about a two terabytes of uncompressed audio aiff. So even if that is compressed it would still occupy a quite large space at a server. (Should maybe have a talk with people at the waag about this.)

DSP (Digital Sound Processing) is possible in CoreAudio through the use of external AudioUnits (AU's). But getting to understand the workings of these AudioUnits on a programmatically level will be most difficult, for I know a little programming, but very little about audio and an external audio specialist might know how to use the AudioUnits in the context of a high level application, like LogicAudio, but probably can not easily translate that back to the code structures underneath.

APPENDICES

notes

Areas of research:

1. search for existing systems
2. interview techniques
3. notation techniques
4. AI (Datamining, CF Agents, Neural Networks, Genetic algorithms)
5. Audio (CoreAudio)

Theoretical research topics:

6. Followers behavior and alternative movement related to search engines and CF. (Or why do I like to walk in fresh snow)
7. Comparison of artificial intelligence techniques in relation to navigating an audio archive
8. Relating an artificial intelligent conversation partner with interview techniques. (How to reach a nice conversation: blathering until you reach some common ground which is more interesting than sharing generalities.)
9. What happens if you let the system wonder off by itself? Will it ever find an interesting thread or combination? The million monkey myth.

Constraints:

time
money
techniques (which I am able to master)
physical body
my own intelligence
comprehensibility (or the intelligence of my audience and users)

List of annotating tools that I currently use:

handwritten:

- index cards
- notebooks and other booklets
- A4 sheets
- moleskin notebook
- post-its (pasted on all the above media)

computer-typed:

- TextEdit (to type notes and use simple RTF functions: e.g. bold, italic)
- SubEthaEdit (to work collaboratively on the same document)
- BBedit (to make sites and quickly publish them)
- wiki sites (to publish notes on a space where everybody interested can co-edit and alter)

spoken notes:

- iPod + iTalk (attached microphone to portable audio harddisk which synchronizes with my music jukebox, iTunes)
- calling my home phone answering machine from my mobile, to leave a message for myself

other:

- taking photo's with my digital camera

research notes

Similar or related projects:

"What was I thinking", a memory prothesis. Sunil Vemuri, MIT Media Lab, <http://web.media.mit.edu/%7Evemuri/wwit/wwit-overview.html>

Ipodder, Podcasting. A personal radio distribution and publishing system which downloads episodes to an iPod, based on RSS. Adam Curry, <http://www.ipodder.org/>

Other interesting links:

Linux in iPods:
<http://ipodlinux.sourceforge.net/>

Possible external sources of knowledge (people):

- people at the "Centrum voor Beeld en Geluid" in Hilversum, where archivers are using computer software to index and annotate historic radio fragments.

- editor, producer of the avonden.
- somebody who can tell me about interview-techniques and technical stuff about making radio programs. (Bente Hamel?)
- programmer or designer that programmed the former 24HCinemaService site of the VPRO, a system that used CF (collaborative filtering) recommendation. Might have practical implementation tips, but also interesting as source of ideas. (But who is this?)
- somebody that could explain me in simple language about artificial intelligent agents. (University of Amsterdam, CWI, Centrum voor Wiskunde en Informatica, faculteit AI, Plantage Muidergracht)
- some biological genetics analysis persons, that could tell me about the current software techniques used in DNA analysis.
- audio engineer (Rens Pluym? Tomasz Kaye? Maarten Schuurman? could design audio environment and might know some useful AudioUnits (AU) that could be used in conjunction with my CoreAudio play backing application.)

lexicon

CVS:

concurrent versions system. CVS is an example of an Open Source version control system. CVS is actually a part of SCM (Source configuration management), a system, which is all about maintaining different versions of the same code or document, so a group of people can all work on a set of documents, without losing recent changes or people working on older deprecated versions.

recombination system:

Or an mutation / recombination system. DNA / RNA as example. A system which uses mutation and recombination techniques (cutting up and pasting back together) to create a desired string of material.

jogwheel / jogdial:

an external human interface device which can be rotated and turned. A potentiometer-like knob. The action of rotation the device is normally used to rewind and forward (scrub) through time-based material (film / audio).

keychord:

a keyboard with just five keys, which can be operated with the five fingers of one hand. Through combinations of pressed keys the user can still type the 26 letters of the alfabet. (5 keys has 32 possible settings, so 31 different characters can be typed, leaving one for not touching the keychord at all).

memetics:

the theoretical and empirical science that studies the replication, spread and evolution of memes (7)

meme:

an information pattern, held in an individual's memory, which is capable of being copied to another individual's memory. (7)

Collaborative filtering (CF) recommendation system:

An artificial intelligent recommendation system that makes use of bots that try to match similar users with similar favorites. First research project was a music recommendation system called ringo.

rhizomatic structure:

a grass-root structure. A non-hierarchical organic linking system that is able to grow new connections from two separate elements and create a vast mesh through these cohesive chains.

references

(1):SubEthaEdit. A collaborative text editing tool that is remarkably coded using itself. <http://www.codingmonkeys.de/subethaedit/>

(2): the cocktail party effect. The phenomenon that a listener is able to cut out high amounts of surrounding sounds and voices to focus on a single voice of the one that he or she is listening to. "The ability to follow what one speaker says in the presence of the chatter of many other is called the 'cocktail party effect'." (John Pierce, "Hearing in Time and Space". In "Music, Cognition, and Computerized Sound. An introduction to Psychoacoustics", Perry R. Cook, MIT Press, Cambridge Massachusetts USA, 2001)

(3): de avonden. the evenings. A daily, two hour, dutch spoken radio program on the MW (mid wave) band, so not the main stream FM band made by the VPRO, a liberal sort-of-left-wing broadcasting organization, with a great member population within the cultural, intellectual, educated elite of the Netherlands. The "avonden" can best be described as cultural magazine, filled with poetry; readings by writers; book, CD, and movie reviews; documentaries in somewhat obscure subjects; and conversations and interviews with artists, architects, writers, philosophers. The program is on 5 days a week and lasts for two hours, without a news and commercial break in the middle of these two hours, which is an exception in dutch horizontally programmed public radio. <http://www.vpro.nl/programma/deavonden/afleveringen/>

(4): Accuracy can be a problem. from Paul Perry quoting an unknown source ("you know who you are"), <http://www.paulperry.net/notes/recommender.asp>

(5): co-designing and co-creation are keywords used in the process of meta-design. "Meta-Design: A Framework for the Future of End-User Development", Gerhard Fischer and Elisa Giaccardi. In "End User Development - Empowering People to Flexibly Employ Advanced Information and Communication Technology", H. Lieberman, F. Paternò, V. Wulf, 2004, Kluwer Academic Publishers, Dordrecht, The Netherlands.

(6): Doors of perception Round Table Conference. (19th november 2004), "What are the ingredients of

success of a design project?" John Thackara, Martine Posthuma de Boer, Kristi van Riet, Dirk van Oosterbosch et al. <http://static.doorsofperception.com/2004roundtable/>

(7): from Principia Cybernetica, <http://pespmc1.vub.ac.be/MEMES.html>