BRAIN STORM

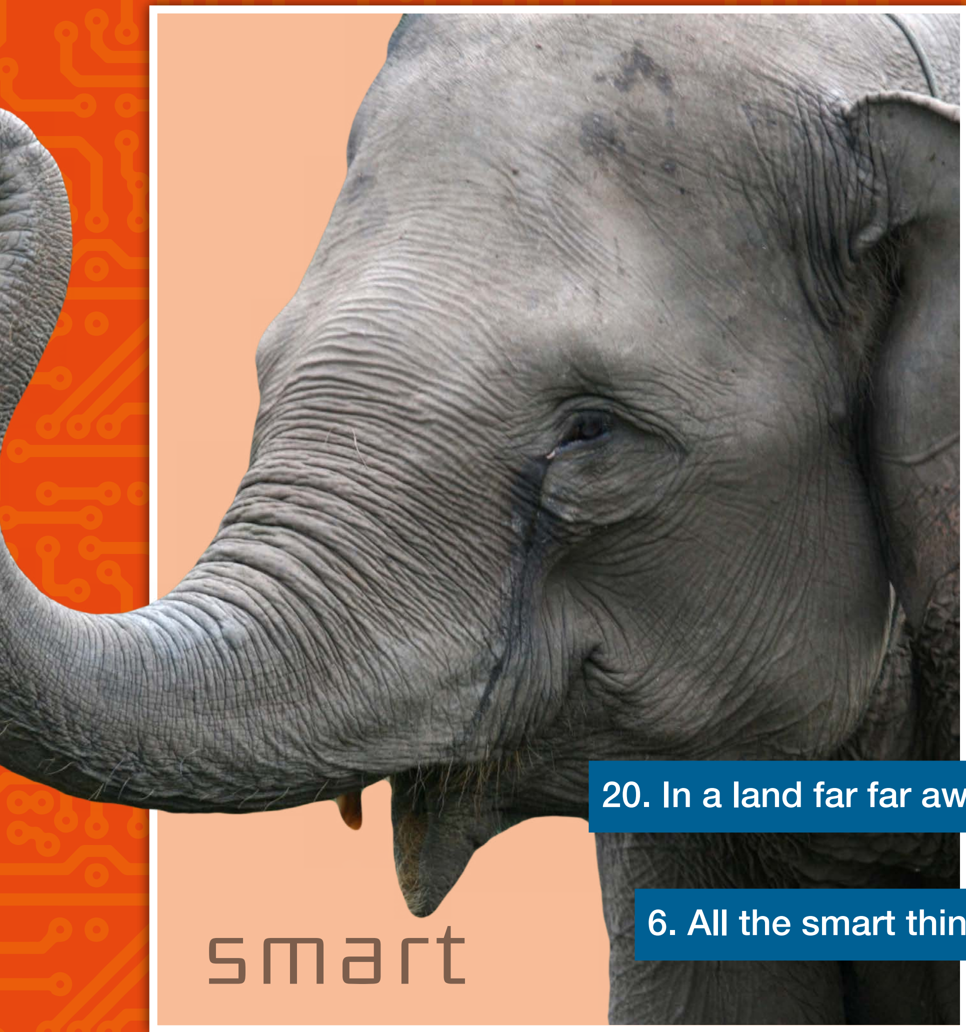

20. In a land far far away

6. All the smart things

smart

TABLE OF CONTENTS

Editorial	3
By the board	4
Comic	5
Column : All the smart things	6
Switch: Eindhoven & Groningen	8
Alumnus AI : Auditory filtering and classification	10
Puzzle	18
MxCee: In a land far far away...	20

Entertaining article, no scientific content. A puzzle or a Cover related article are examples of this category.

Easily readable article on a scientific topic. Should be comprehensible, even without any prior knowledge.

Scientific article that explores a certain topic in depth. Might assume the reader has taken a course that's related to the topic.

EDITORIAL

BY: Sanne Bouwmeester, editor

The theme is of this Brainstorm is smart. Smart is a term you hear often. This person is smart, that person is not. Are you smart enough, or not? When thinking about smart, most people think about reasoning, learning ability and memory. People often try to determine intelligence through IQ, but EI, your ability to recognize emotion and anticipate those of others, should not be forgotten either. There is even something called the Flynn effect which entails that people are constantly getting smarter. You can also become smarter by reading this Brainstorm!

If you are not among the smart people, no need to feel distressed! Nowadays, there are a lot of smart devices, such as your Smartphone, smart TV and your smart watch. They can help you with just about anything. Planning, cooking, trips and vacations, taking pictures – the results can be seen on the photopages – and work of course. All these smart applications make our lives a lot easier.

However as smart as these applications may be, they still can't make magazines for us. That is what the devoted members of the Brainstorm are for. As a junior member I am learning a lot and becoming smarter every day. As a candidate board member I am learning a lot as well. We learn all our live, even the 80 year old grandma's are.

There are people who were clever enough to sign up for the MxCee trip. They travelled across the world and became even smarter than they already were. You can learn from their intelligence and read about their adventures in New Zealand in this Brainstorm. Some of the participants were part of the Brainstorm. We combined their experience with our exsting knowledge to make this Brainstorm smart. Continue reading to get a glimpse of it.

If you feel smart enough already you can try the puzzle. You can test your logic skills and try to derive who lives where. Enjoy your Brainstorm!

PHOTO: Sanne Bouwmeester

BY THE BOARD

BY: Guus Klinkenberg, treasurer

Nowadays, a lot of things are considered to be smart. It varies from individual devices, such as phones and boards used in classrooms, to entire principles, such as grids. Smart systems are often considered smart when they have the ability to analyze the current situation and react to it. When thinking about smart systems using that definition, there are many opportunities for smart systems. Even more is possible when systems are interconnected.

Smart is also an acronym in project managing and plays an important role in SCRUM development, where it is used to describe Specific, Measurable, Attainable, Relevant and Time-bound goals. We will come into contact with most of these things during our study, but the one thing that I want to focus on, is S.M.A.R.T..

S.M.A.R.T. is a system that has been developed for storage devices. It tries to detect errors based on various sensors and tests. When it thinks that the drive might lose access to data soon (for example because various tests are failing), it will inform the user of the imminent failure. It is up to the user to actually do something with it, but it is a very nice system and has saved me from data loss various times.

A similar construction can be seen within our beautiful association. The board sees and hears a lot of different things, and when it thinks that something will go wrong, it will report it to the General Assembly. There is one big difference, being that the board is able to fix things itself, while S.M.A.R.T. really only is an observatory system. This is a good thing, because we would not want to call in a General Assembly for every little thing that needs fixing, because then we would have a General Assembly every week, if not multiple times a week.

Anyway, the board receives a lot of information from a lot of sources, sensors if you will, and processes this when making decisions. Using this information we try to set goals that are very smart and therefore I think we may call ourselves part of a smart system.

PHOTO: Guus Klinkenberg

COMIC

BY: *Joke Kalter*

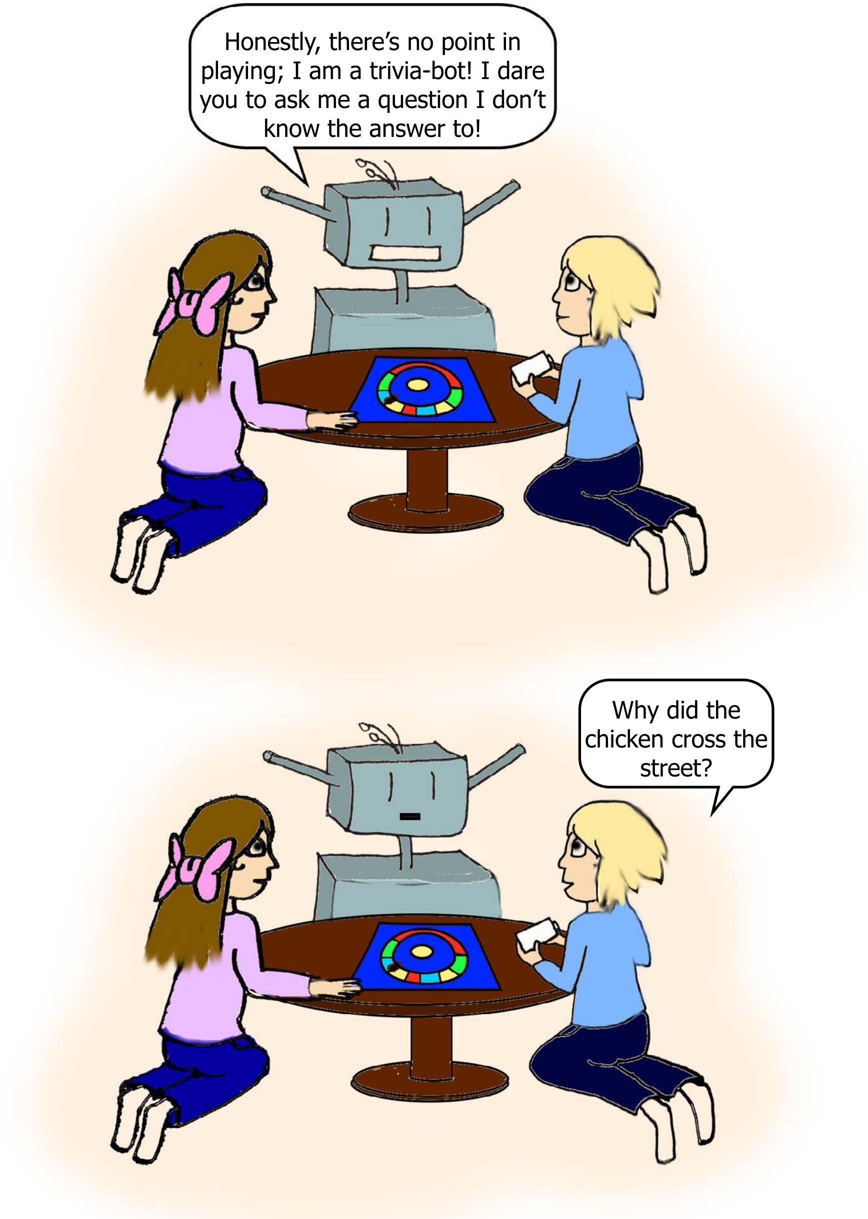

ALL THE SMART THINGS COLUMN

BY: *Arryon Tijsma*

The Brainstorm asked me to write a column about 'smart'. Besides it being a most suited subject for my person, I feared some of you would not understand the concept. Let me elaborate. This is what the Oxford dictionary calls 'smart':

> intelligent, or able to think quickly or cleverly in difficult situations

Not that this definition really has anything to do with the rest of the column, but I just wrote 86 words that you've still read, and a quote looks nice on the page. Smart, right?

Now, if you've given up on life because 'smart' does not apply to you (if you're not studying at the FMNS faculty for instance): fear not. Because using my crystal ball I have seen the future of smart things and using this column I will make your life worthwhile again (yes, I'm really that smug).

Because the near future will unfold many inventions that will outsource the need to be smart entirely. Don't give up hope, in the future everything will be good and easy! I predict that within ten years even the most retarded person will be able to live an impressively smart life. If you haven't guessed the reason yet, it is, naturally, the *internet of things*, a.k.a. *ALL the smart things*.

The internet of things is a development where every device gradually becomes smart. Soon, every device will have a sixty-four core integrated processor so you can mine bitcoins when brushing your teeth. I present to you: an overview. Even with images, if you're struggling to keep up!

The smart beer glass

A smart beer glass which sends a Facebook friend request when you clink with someone else already exists, but that's just child's play. Beer glasses will get so much smarter. For instance, a glass from the future will indicate the alcohol percentage of your drink, so you never have to accept watered down beer anymore. It will also show the rate at which you're necking your beer, so you can brag to your bros that you're really the fastest, or get comatose trying. Naturally, to unlock all your beer glass features, such as aggregated statistics and challenges and achievements on social media you will have to get a subscription and pay a monthly fee. But that's just progress.

The smart toothbrush

Oh nevermind, that one already exists. Really? Yes, really [1]. It doesn't mine bitcoins, but you can play games while brushing your teeth. Maybe you can use it as a controller for your Wii and play Smash Bros. While brushing. Dentists will go extinct in the future.

The smart bicycle

While cycling, you can track your kilometers and your route, but that's all done using a smartphone. In the future, bikes themselves will come in 'smart' editions. They can do a whole range of things. When you are late for the train or bus home to your parents again and throwing down your bike somewhere around central station, instead of just lying there rotting in the rain being stupid, your bike will get up again and cycle itself to the nearest free parking spot. Maybe even more convenient will be the recall function. After a decent evening going out, people will shout out in a slurred manner: *"where the f*ck did I leave my bike?!"*. Their smartphone will pick up the distress signal and voilà: the bicycle will autonomously approach its owner.

Smart smartphones

The future will have smart smartphones, just like we have homo sapiens sapiens. Because of the internet of things, smartphones will move up a rank in the digital hierarchy. Smartphones will have smart devices of their own (yo dawg...), delegating their tasks. You will have to feed your smartphone in order to keep it happy, and walk around with it enough. In fact, smart smartphones will outsmart us by being connected to the internet of things. They will have their own wearable accessories like a young boy studying Artificial Intelligence. Does that scare you?

Remember that smart smartphones will essentially be able to do everything for you, so you don't have to anymore. You can just rack up statistics on your beer glass while playing a skiing game with your toothbrush, waiting on your bicycle to pull up. What a good life :)

Of course I'm just scratching the surface and giving you a teaser of everything the future has in store for you. Now that you feel invigorated thanks to the prospects of a better world, you can continue reading the other articles in this edition. If you still don't believe me, the following picture will convince everyone with irrefutable proof:

[1] http://www.kolibree.com/en/

DIAGRAM: *The proof*

SWITCH: EINDHOVEN &

BY: *Martijn Luinstra & Bor de Kock*

1

It is early in the morning after the Cover borrel when I wake up to take a train to Eindhoven, a journey to the other side of our country that one does not take without a good reason. I am on the quest to explore the city in the land of carnival, the country's capital of technology.

My first stop is the room of GEWIS on the university campus, which is just a five minute walk away from the train station. This is where I wait for my host and tour guide, Bor, who currently is in a meeting. After a while he arrives and we leave the university for a walk through the city. The city of Eindhoven in its current form has a relatively short history that has been influenced heavily by its university and the company Philips, making it the Silicon Valley of Western Europe. Many old Philips factories are turned into houses or cultural platforms to serve an above averagely intelligent community and its surrounding rural area.

Our tour ends at the Philips museum, where we are greeted by the museum's oldest employee who quickly shows us the museum's oldest document. The museum itself is a propaganda exhibition about the history of Philips, showing the company's history from the development of the first light bulb to the modern LED. Besides showing off, the museum guides us through the absurd history of electric shavers and the wonderful design of their 1960s advertisements.

After having a brief tour through the old city centre, it is time to explore something even more interesting: the city's student life. Back at the university's campus, we visit the weekly borrel of GEWIS, which was organized by their dispuut for outside activities this time, and therefore takes place in the freezing cold of the Dutch winter. To make the cold bearable, the dispuut provided campfires, freshly barbecued hamburgers, snert, glühwein and beer.

To end the day in style, we visited the study association of Electrical Engineering (e.t.s.v. Thor), which held a party to say goodbye to their bar in the catacombs of a university building, which they needed to leave after decades. I made my way through a partying crowd into a sweaty, overcrowded bar which they call "The Walhalla". After having had way too many drinks and having met some amazing people, it was time to travel back home. While writing down some notes, I know for sure that this won't be my last trip to Eindhoven.

PHOTO: *Martijn, Bor and others in Eindhoven*

GRONINGEN

After Martijn's visit to the sick sunny south in February, it was time to do some sightseeing in that shaky city up north.

There is nothing above Groningen, is what they usually say, even though we managed to elevate ourselves some 311 steps more within the Martinitoren to enjoy the views of this city and the views beyond its borders. Seeing Groningen from above helps you understand why it's called 'Stad' in their local dialect: where the suburbs of Groningen end, there is only rural land left.

Apart from climbing towers, visiting the University Museum, strolling through the Prinsentuin, and enjoying the grass (the grass!) of the Noorderplantsoen, we of course visited that huge yellow-and-blue tourist attraction the north has to offer: the largest IKEA of our country. Martijn and his friends had their Sweden Study Trip Reunion Dinner, and I bought some furniture to take home, because why not. Also, I had never seen a shopping cart escalator before, how epic is that?!

As I am a Study Association enthusiast myself, we visited the Cover Room to enjoy some coffee, complaints about the General Assembly, a stressed out study trip committee and a lot of people playing on the Wii. In a lot of ways Cover is apparently not that different from GEWIS, it's just that we do not have Maikel sitting on things.

One thing the Groningers seem to take a lot

PHOTO: *Brakke selfie in Groningen*

of pride in is that the bars are open late, and a typical borrel starts at 21:30 (so people show up around 23:00), and does not end until the sun rises. In Eindhoven the bars close at 2:00, but since our borrel starts at 16:30, the effect is usually kind of similar and we even get to sleep semi-properly afterwards.

The final part of the program was of course Cover's candidate board announcement borrel, where I thoroughly enjoyed the new Cover Server Door Logo: having dragged that hellishly heavy thing from Eindhoven to Groningen last year, it's good to see that they are actually making good use of it. The four kandies that were able to stand upright and talk seemed really nice as well.

And then, after a lot of drinks, water pistols, almost-brassed kandiedassen and a few hours of sleep, I made my way from Groningen, through Drenthe, Overijssel, Gelderland, Utrecht and again Gelderland, back to Brabant. We'll meet again, 050!

BY: *Arryon Tijsma*

Environment aware auditory filtering and classification

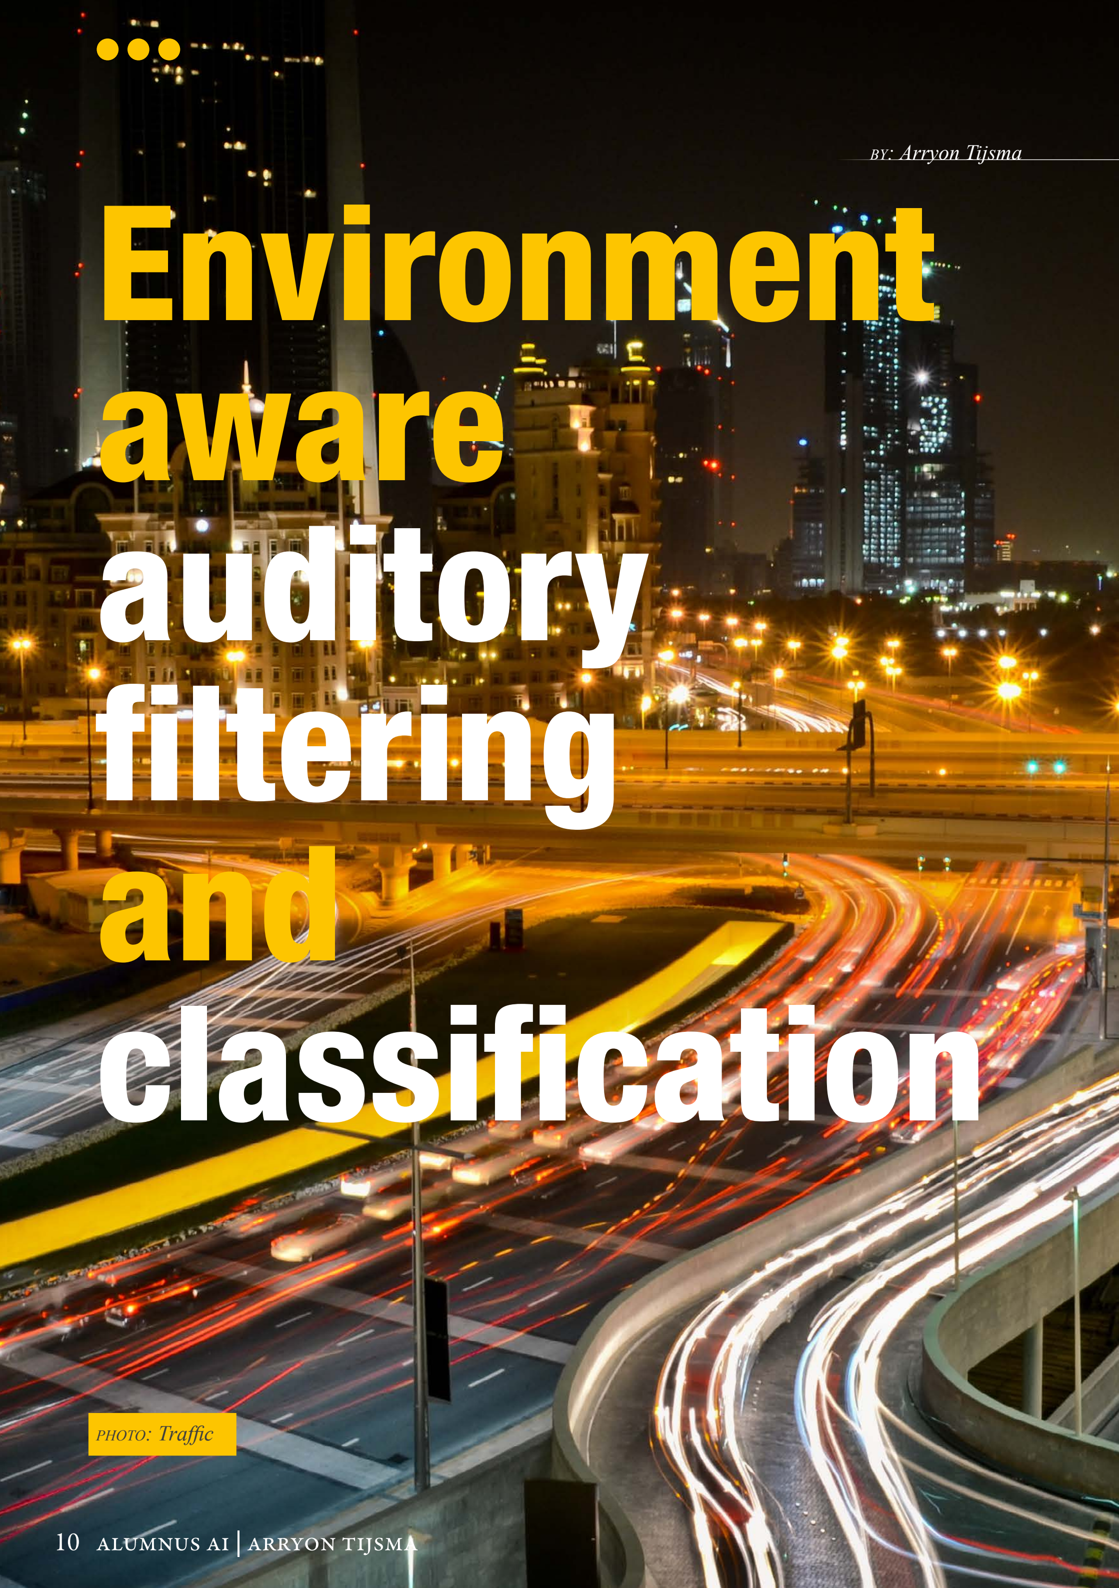

PHOTO: *Traffic*

Abstract

In this article we treat the concept of everyday listening to audio environments, by designing a machine learning classifier, using Gaussian Mixture Models, that learns to classify the audio environment based on naturally occurring texture features that have been shown to be present in different types of signals. We show that it is possible to classify different audio environments with an accuracy of 85.8%, and that such a system can also recognize possible speech-interfering surroundings. We also present a local saliency adaptive filter called the Focus-On-Interest filter, which is capable of filtering these speech-interfering components.

Introduction

Listening to an environment in the same way a human does is a trait not many artificially intelligent systems have focused on. One can argue that it is because the right ways in which to listen have not been used, in order to easily distinguish between different sound producing sources. Surely, when we hear what is around us, it seems almost trivial to determine the source of a sound. Rather than listening being easy, simultaneous processes make it seem easy for us.

First, our ears are perfectly fine-tuned for sensitive perception [6, 7]. Second, there are auditory mechanisms that can block out unwanted sounds and focus on the part we want to recognize [21, 12]. Third, we have a vast library of knowledge to compare our perception to: our memory. The memory can prove to be very efficient at recognizing the correct sounds [22], so much that it sometimes fools us into believing we hear things that are not there [20].

The way in which humans hear has been dubbed everyday listening by Gaver [13]. It is not the focused listening we do when trying to appreciate music.

When you walk through the streets of a city, it is likely that you will not pay attention to how a loud mechanic rumbling sound sounds, rather than what it is that produces that sound – likely to be a large car – and whether the object that produces that sound is worthy of extra attention. This specific scene is an example of everyday listening.

Everyday listening is the almost subconscious awareness that can alert us, and help us move through the world. It is the kind of listening that can suddenly alert you, and possesses an always-on reaction system that is almost as old as life itself [10, 19]. Much auditory research uses signal features of audio that are designed to work best for speechrecognition and classification: so called Mel Cepstrum Cepstral Coefficients (MFCC's) [14, 18].

We can read, mainly by Gaver [13], but also by other scientists, that when sound is produced, this is always because of a physical process in the world: either a type of friction, or something slamming, or a repetitive motion, or something else entirely, is the creator of the sound [15, 5, 3]. This means that from this origin, characteristics are interwoven in the audio signal, much like the texture of a piece of fabric [9]. We should therefore be looking at such features to classify audio sources using machine learning, much more than using features that are tailored to the domain of speech recognition.

In this article we propose using a type of texture based audio feature to extract features from audio and classify the auditory environment based on their occurrence. For the actual classification, a Gaussian Mixture Model (GMM) is used clustering the audio into the multidimensional feature space using probabilities. Besides a classification method, another mechanism that makes up the technique of everyday listening is developed: a temporal filter based on the rise and

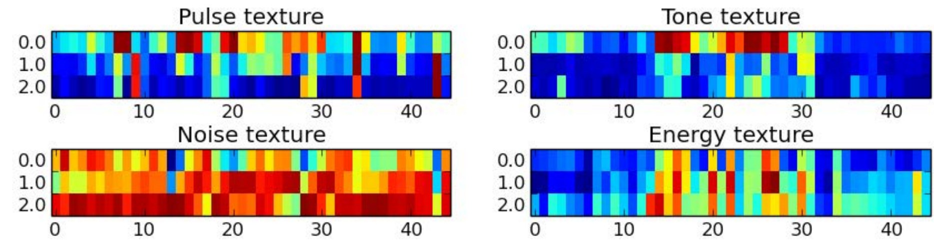

Figure 1: Texture matrices after the final feature extraction stage. This data is vertically stacked and fed to the GMM classier

fall of signal intensity that accomplishes focusing on the foreground of a signal, and is dubbed the Focus-On-Interest (FOI) filter.

Audio algorithms

Texture features

The idea that sound can have a texture-like quality much like the structure of cloth, because every sound is a consequence of physical interactions of real-world objects, was first introduced by [2] as property of sound to be used for recognition. There may exist similarities between similarities in sound, and similarities in structural patterns in the time-frequency domain of sound, hence we refer to this concept as a texture [9].

As stated before, this concept can be linked directly to earlier research from Gaver. Gaver specifies three separate physical causes for a sound: impact, resonance, and turbulence. An impact sound produces a time-localized, vertical time-frequency intensity spike. A resonance sound produces a frequency-localized, horizontal time-frequency intensity spike. A turbulent sound produces a pattern that is neither of the previous, and is unstructured in both time and frequency [13].

Van Elburg et al [8] have formalized how texture information can be described, resulting in features defined in the time-frequency domain called *tract features*: $\mathbf{T}_{-}$ and $\mathbf{T}_{|}$ (horizontal texture and vertical texture), that have shown to possess distinction probabilities that significantly differ from randomly structured sound [9]. Conversely, from $\mathbf{T}_{-}$ and $\mathbf{T}_{|}$ a third structure, $\mathbf{T}_{n}$, can be derived, which correlates with the other two and represents unstructuredness in the sound. All features as presented have a masking property. To make them usable for thresholding, they need to be transformed into a soft-margin variant using a threshold θ.

Before the texture features can be used in a classification algorithm, they need one final processing step: applying them to the actual time-frequency signal $\mathbf{E}$, which yields energy-related values instead of masking (boolean) values. This, in turn, makes the features suitable for many machine learning applications working in continuous domains, like the GMM we use in the experiments.

Temporal leaky integrator filter

One theory on the use of focusing mechanisms when hearing states is that we have many filters in the part of our brain that processes sound with different influences based on the temporal development of the sound, and the brain can switch to a different filter when we want, whether conscious or unconscious, for instance in case we could not clearly pinpoint a sound [17].

A temporal filter developed by Andringa accomplishes this. The core idea of the filter is to vary a time constant, τ, by leaky integration as a function of the local ratio of foreground to background of a signal in the energy domain.

This means a fourier transform or cochlear model must already be applied to a signal before using the filter. In our experiments we used a cochlear model to create a frequency/time distribution of the energy of a signal, called E.

Leaky integration must be seen as the type of summation over time, like biologically inspired neurons do, with a decay (loss) component λ^{-1} applied on the energy of the current time step, E_{t-1}. For ease of concept, we will refer to foreground when we mean the unaltered energy representation of the signal. The background BG is the leaky integrated energy of the previous timestep, according to

$$\mathrm{BG}_t = \mathrm{BG}_{t-1} \times \lambda + \mathrm{E}_t \times (1 - \lambda)$$

and is a value that always converges towards the foreground. How fast this happens is dependent on the loss, λ. It is dependent on the ratio between foreground (unaltered energy) and background (lossy integrated energy), the Foreground-Background-Ratio (FBR). When the *FBR* is negative, this means the unaltered energy has higher intensity than our supposed background. In that case, λ will grow smaller rapidly and the background will converge faster towards the foreground. When the *FBR* is positive, λ will grow larger slowly, and the next value of the background will converge slowly towards the foreground.

How λ changes thus shapes how our filter works. In the algorithm, λ is changed by, at each time frame t, calculating the *FBR* by subtracting BG_t from E_t. The correct value of the loss is then found with a lookup in a pre-calculated loss vector $\mathbf{\Lambda}$. Looking up which value to choose as loss λ creates a dependency to the local shape of the signal. Specifically, to the local discrepancy between energy and background component, as we will see later. This has the benefit that filtering is done with adaptation to how salient the sound is at the current time.

Methods

Sounds

Audio data was downloaded from the Freesound.org public database [11]. Only WAV audio was selected for the experiment. Any stereo tracks were saved as mono by selecting only the left channel. We made sure all audio was saved as 16 bit PCM audio with a sample rate of 44.1kHz. In total 30 minutes of audio was selected by hand from four audio environment classes deemed interesting. The classes are: **traffic**, **natural**, **indoor**, **chaotic**. Sounds were selected so they balanced the total duration between all classes.

The choice for these classes is inspired by research on affect [1, 16], stating that the affect brought on by sounds can be divided into four quadrants: moderate human sounds for lively, natural sounds for calm, noisy and homogeneous trac sounds for passive (maybe even hindering), and a very loud cacophony of sounds for chaotic.

Besides the sounds from environmental classes, the same audio files were also used to classify based on speech hindering components in the sound. To determine what a hindering component is, the same speech fragment was pasted onto all environmental recordings, and made sure to be clearly audible with its mean intensity lying 9dB higher than the mean intensity of the background fragment. We then made a selection based on the criterion "*would this recording hinder me as a human to clearly understand the speech that is done?*". The possible answers are either "yes" or "no", so the resulting two classes to classify with a machine learning classifier.

Feature extraction and classification

Before we start training a classifier to recognize the mentioned four classes, we need features to do so. The algorithms used for extracting the mean $E_{\_}$, $E_{|}$, and $E_{\mathbf{n}}$ are suited for this task. We configured the algorithms so that for each second of audio, transformed into a 133 × 200 matrix, the output of step 4 in our software process is a 4 × 10 matrix per feature. Due to merging in the tract features, not all cochleogram bands can be used, discarding the first 10 and last 10 bands. To further simplify division, we discard another three bands. The frequency splits occur in bands 10–36, 37–63, 64–96, and 97–118.

Since all feature values are in the same domain, no extra transformation like PCA analysis or z-scoring needs to take place. If we consider the output matrix from the texture extractors, we directly have input features for a machine learning classifier. The resulting output for the three features $E_\_$, $E_|$, $E_{\mathbf{n}}$, and an additional, down-sampled version of **E** to the same format, is a matrix of 12 dimensions and a length equal to the sample rate times the total number of seconds of audio in our entire data set. An example can be seen in figure 1.

We are being purposefully very course in splitting up the frequency bands, because classification should not be about exactly where in the frequency spectrum a sound happens, but rather how much the characteristics of certain broad-band frequencies are occurring in different auditory environments, and if this information can be used to accurately separate these environments.

For all sound fragments, we calculated their texture extraction responses, and saved the result in one data set with labels. In other words, we treat the column vectors (data points) in our sound features in the same way as a *bag of words* is treated: separated from their context. All in all, we obtained roughly 18.000 data points.

Upon inspecting a 2D plot, it became clear that many data points were placed in roughly the same position in the multidimensional space. We can see this very clearly if we plot only pulsatility

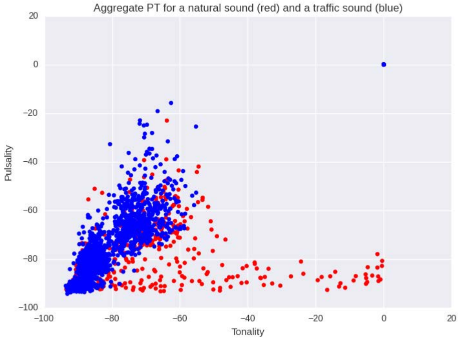

Figure 2: Projected 2D plot of 4D features E$_|$ and E_

versus tonality, two features that according to a PCA analysis contributed most to the variance of the data set. In figure 2 we project all four frequency bands of $E_\_$ and $E_|$ onto the same axis, and see that many data points actually denote a large absence of both features. Listening back to the audio, this is easily clarified: those data points are simply reflecting the moments where nothing happens and only microphone noise is audible.

Since the characteristics of a digital recording are probably roughly the same, we can expect much overlap from all four classes in this area during moments in the recordings where nothing or almost nothing is audible.

Data was trained using a Gaussian Mixture Model classifier, using the Expectation-Maximization (EM) algorithm to estimate the sample parameters $\widehat{\mu}$ and $\widehat{\sigma}$. Since we have labels, we already know the mean of the data points belonging to one class, $\{x_j \in C_j\}$. Only the covariance matrices have to be estimated by the EM algorithm. Our choice herein is that we can specify four shapes of the covariance matrix: *spherical*, *diagonal*, *tied*, and *full*. Spherical entails that every covariance dimension is kept equal per model, resulting in a spherical boundary. Diagonal restricts to a diagonal covariance matrix, because it is both easier to compute, and might generalize better. A tied shape is the sharing of

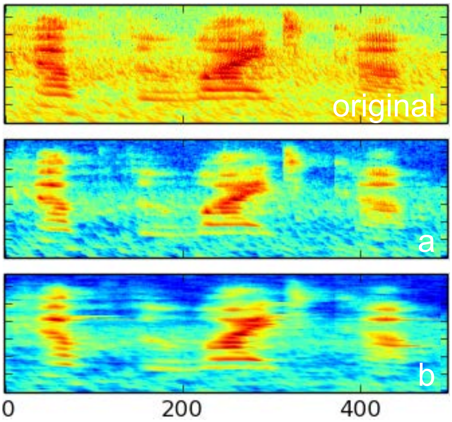

Figure 3: Two models of the same cochleogram: (a) fast 2.5s and (b) slow 30.0s τ response

equal variables (covariance matrices) between multiple models.This means the models will have the same shape in our feature space. Finally, full covariance is the unrestricted version of a GMM that does not restrict anything.

Other parameters of the gaussian mixture model algorithm as provided by Python Scikit are kept constant: 1 initialization (we do not use random state), and 20 iterations of the EM algorithm to converge.

Training the classifier is done by applying five-fold cross-validation on the data set with a 90%/10% training-test split. For each fold, and for each covariance type in each fold, model parameters are estimated using EM, after which testing is performed. The F1-score of accuracy on the test set is calculated, and later averaged over folds. The averaged accuracy is the performance estimator of the classier.

Focus-On-Interest filter

To create the FOI filter, two leaky integrator models were used. Their sets of parameters does notn't tell us directly the way they transform a signal, so figure 3 gives an illustration of a processed sound fragment that was used to develop the filter. It depicts the sound of traffic on a highway, in the intermediate vicinity, and a voice command: "Sudo, pick up the red cup". Two time constants τ were used: one of 2.5 seconds, and one of 30 seconds.

Notable is the smearing effect that occurs when we apply the filter with a large time constant τ to **E**. This occurs because the calculated background lags behind the real signal and is in some ways a running average with sudden onset decay when decreasing. This also shows because energy is shifted slightly towards the right.

Figure 4a shows the resulting absolute difference between the two models of different time constants τ. As we see, if we choose the correct short and long constants like in this example, surrounding a rise and fall of intensity there appears a region with an absence of energy. In the regions where energy is constant, on the other hand, more energy stays present. This means there is not much difference when filtering for fast changing sounds, then when filtering for slow changing sounds, and is thus an indication of a non-salient signal.

Figure 4b shows what is left of the energy when we subtract the difference of the two models from the unaltered **E** matrix, and threshold it to highlight the difference. Most of the speech signal remains intact in the onset of the signal, and at the offset

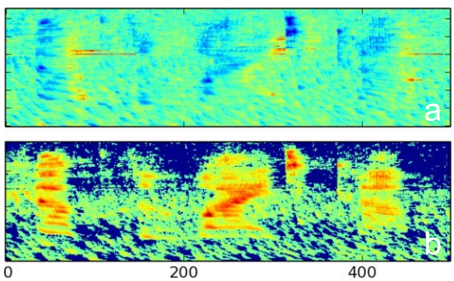

Figure 4: Filtering based on saliency: (a) difference between two models, (b) thresholded substraction

not much is subtracted because there is not much difference between the two models.

Results

For the FOI filter, no real results were obtained in terms of improvement of accuracy in speech recognition or classification. The only results we can show look promising: because the FOI filter adapts to local changes in intensity, based on two time constants, we can suppress a longer lasting energy component with less intensity (ambient noise), keeping intact the shape of the salient foreground (the speech signal). For instance, when a lot of noise in the background is audible, but nothing apart from that happens, the FOI filter will throw away almost nothing of the signal: the noise contains the salient information in this case the salient information. But as soon as a more intense sound becomes audible on top of the noisy signal, the local adaptation will emphasize the current salient foreground (because it is a rising intensity) while discarding most of the noise in the vicinity of that saliency (because of the smearing effect which creates a margin). As soon as the extra sound fades away again, background will not be thrown away and will become salient again. This whole process much resembles the focusing mechanism as described by [17].

Our GMM classifier can of course be tested on accuracy. Based on the five-fold cross-validation, the outcome of our covariance schemes is that the *diagonal* covariance performs best on our data set. It is sometimes suggested that the tied or diagonal schemes are only successful when the data set is insufficiently large, so that may be more of a symptom of the data rather than the true best method [4].

When looking at classification of environment class, the average F1-score with *diagonal* covariance is 85.8%. The confusion matrix is depicted in table 1. When looking at classification of the presence of a hindering component, the winning covariance type is *tied* the average F1-score with tied covariance is 82.3%.

	nature	**people**	**traffic**	**chaos**
nature	99.6%	0.4%	0%	0%
people	0%	96.9%	3.1%	0%
traffic	0%	0.5%	99.4%	0.1%
chaos	0.5%	0.6%	17.8%	81.1%

Table 1: Confusion matrix between classified data points

Conclusion

We manually gathered about 30 minutes of audio from four different auditory environments. Considering no selection took place on what was actually in the recordings, but rather on the question if they were representative of their class, and considering the results of our classifier, it is clear the texture features have some form of salient information in them, enough to recognize with high accuracy the four different environments.

The resulting confusion matrix for one of the folds, depicted in table 1, shows that most confusion happens between *chaotic* and *traffic* sounds. Listening to their recordings seems to confirm this confusion, because there is much audible unstructuredness even in the loud parts of the signal. The same goes for recognizing a potentially interfering sound environment, a task on which the classifier scores almost equally high.

Discussion

A few remarks can be made about the methods used in the article. First, every data set of such small size is inherently flawed. More data should be used in order to test the generalization ability of auditory texture features. Second, the used classes are, even though scientifically supported, arbitrary in the sense that they have not been extensively researched and tested in different combinations. More research about how the features behave in a different classification space should be done in order to appreciate the extent in which they separate different auditory classes. Third, the absence of sound was not taken into account in the research. A Gaussian Mixture Model will always assume that a data point belongs to a certain cluster, but in every recording silence can occur. In silence, as we saw, almost no salient information is present. Therefore, silence should either be taken into account as belonging to its own cluster, or be removed from the data points before classification. Since that in turn creates new problems (when is there silence?), it is not a trivial hurdle in recognizing an auditory environment.

References

[1] Tjeerd C. Andringa and Kirsten A Van Den Bosch. Core affect and soundscape assessment: fore-and background soundscape design for quality of life. Human Evolution, 3(1):19-43, 2013.

[2] Tjeerd C. Andringa and Renante Violanda. The texture of natural sounds. Journal of the Acoustical Society of America, 123(5):3413, 2008.

[3] James A. Ballas and James H. Howard. Interpreting the language of environmental sounds. Environment and Behavior, 19(1):91-114, 1987.

[4] Je A. Bilmes. A gentle tutorial of the em algorithm and its application to parameter estimation for gaussian mixture and hidden markov models. ICSI, 4(510):126, 1998.

[5] Sean W. Coward and Catherine J. Stevens. Extracting meaning from sound: Nomic mappings, everyday listening, and perceiving object size from frequency. The Psychological Record, 54(3) 2011.

[6] Peter Dallos. Cochlear amplication, outer hair cells and prestin. Current opinion in neurobiology, 18(4):370-376, 2008.

[7] Engbert de Boer. Auditory physics. physical principles in hearing theory. i. Physics reports, 62(2):87-174, 1980.

[8] R. A. J. Van Elburg. Signicance in spatial temporal structures in time-scale representations, 2012.

[9] R. A. J. Van Elburg. Structure in sound: Texture analysis, 2013.

[10] Richard R. Fay and Arthur N. Popper. Evolution of hearing in vertebrates: the inner ears and processing. Hearing research, 149(1):1-10, 2000.

[11] Freesound.org. Freesound: collaborative database of creative-commons licensed sound for musicians and sound lovers, 2015.

[12] Jonathan B. Fritz, Mounya Elhilali, Stephen V. David, and Shihab A. Shamma. Auditory attention focusing the searchlight on sound. Current opinion in neurobiology, 17(4):437-455, 2007.

[13] William W. Gaver. What in the world do we hear?: An ecological approach to auditory event perception. Ecological psychology, 5(1):1-29, 1993.

[14] Tomi Kinnunen and Haizhou Li. An overview of text-independent speaker recognition: From features to supervectors. Speech Communication, 52(1):12-40, 2010.

[15] Roberta L. Klatzky, Dinesh K. Pai, and Eric P. Krotkov. Perception of material from contact sounds. Presence: Teleoperators and Virtual Environments, 9(4):399-410, 2000.

[16] Peter Kuppens, Dominique Champagne, and Francis Tuerlinckx. The dynamic interplay between appraisal and core aect in daily life. Frontiers in psychology, 3, 2012.

[17] Christopher A. Kurby and Jerey M. Zacks. Segmentation in the perception and memory of events. Trends in cognitive sciences, 12(2):72-79, 2008.

[18] Robert G. Malkin and Alex Waibel. Classifying user environment for mobile applications using linear autoencoding of ambient audio. Proceedings. (ICASSP'05). IEEE International Conference on, volume 5, pages 509- 512 Vol. 5. IEEE, 2005.

[19] GA Manley and R. Ladher. Phylogeny and evolution of ciliated mechanoreceptor cells. The senses: a comprehensive reference, 3:1-34, 2008.

[20] David G. Payne, Claude J. Elie, Jason M. Blackwell, and Jerey S. Neuschatz. Memory illusions: Recalling, recognizing, and recollecting events that never occurred. Journal of Memory and Language, 35(2) pages 261-285, 1996.

[21] Josef P. Rauschecker, Amber M. Leaver, and Mark Māhlau. Tuning out the noise: limbicauditory interactions in tinnitus. Neuron, 66(6):819-826, 2010.

[22] Norman M. Weinberger. Auditory associative memory and representational plasticity in the primary auditory cortex. Hearing research, 229(1):54-68, 2007.

By: *Steven Warmelink*

PUZZLE

Five Cover members decide to all move into a new flat. However, they all live on a different floors, ranging from the first to the fifth floor. Each member has a unique name, favorite color, beverage, instrument and preferred programming language.

Can you determine whose preferred programming language is Scala? Then send the person's name to brainstorm@svcover.nl before july 16th to have a chance to win a prize!

1) Bastiaan's favorite color is red.
2) Emilio writes code in Haskell.
3) Davey drinks wine.
4) The person who plays the trumpet has a person living below or above him who drinks tea.
5) The person whose favorite color is black drinks rum.
6) The person who plays the guitar codes in Python.
7) The person whose favorite color is white plays the piano.
8) The person living on the third floor drinks beer.
9) Cedric lives on the first floor.
10) The person who plays the trumpet lives above or below the person who codes in C.
11) The person who codes in Java lives above or below the person who plays the piano.
12) The person who plays the Double Bass drinks coffee.
13) Aliene plays the Harp.
14) Cedric lives above or below the person whose favorite color is orange.
15) The person whose favorite color is black lives right below the person whose favorite color is purple.

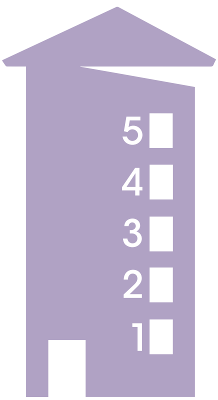

programming language	color	beverage	name	instrument

{kxa} De dataspecialisten

Software Innovations

KxA software innovations is gevestigd in de provincie Groningen. Het is een uniek bedrijf dat innovatieve, gekke, grote, kleine, spannende, mooie, maar natuurlijk ook normale maatwerk software-opdrachten uitvoert. De overeenkomst tussen al deze projecten is dat het gaat om data in alle vormen en maten, bijvoorbeeld:

- Het Nederlandse verkeer in 400 miljard metingen toegankelijk opslaan
- In een stal het gedrag van koeien monitoren
- Software ontwikkelen voor de gigantische SKA radiotelescoop

van BIG BANG tot

- Hulpverleners in de zorg ondersteunen met VCA

Werken bij KxA

Bij ons vind je allerlei achtergronden (natuurkunde, informatica, AI, etc). Iedereen deelt het enthousiasme voor softwaretechniek en wat je daar allemaal mee kunt doen.

We hebben regelmatig afstudeeropdrachten, stageplekken én vacatures. Je krijgt hierbij een opleidingstraject om je helemaal in ons vakgebied te bekwamen.

Ben jij geïnteresseerd in het werken bij een High Tech bedrijf? Kijk dan eens op **www.kxa.nl**, of neem contact met ons op via mulder@kxa.nl

BY: *Jonathan Hogervorst, Maikel Grobe, Jonathan Maas and Annet Onnes*

MxCee: In a land far far away...

PHOTO: *Auckland*

Jonathan Hogervorst:
Back in October, twenty students decided to make a trip to New Zealand. We couldn't depart right away: to earn a flight ticket, we had to work on a case for Crowdynews. Divided into different groups, we built a system to generate news articles from Twitter trends. After a hundred hours we completed the project and were allowed to board the plane.

We arrived in Auckland, where we stayed in a hostel for a few days. During daytime we walked through the city and enjoyed the beautiful view from the summit of Mount Eden, which was once a volcano. Another day we visited Waiheke Island. Thursday night we went on a pub crawl and joined a poker tournament in the hostel. Lotte managed to beat one of the few non-Cover participants in a nerve-wracking finale. Due to the high beer prices in New Zealand, the prize money didn't last very long.

We also visited the Auckland University of Technology. After a lecture about some familiar topic, we met the epic Professor Wai Kiang 'Albert' Yeap. He enjoined us his theory on AI: rather than building systems with perfect performance, we should try to build systems that mimic human cognition.

Some days later we visited the first company, Rush Digital. They forgot to prepare a presentation, but had arranged drinks. It is still unknown what the company actually does.

After leaving the hostel, we went to pick up our motorhomes. Dividing twenty people over four campers is easier said than done: it took the MxCee much effort to construct a habitable division. Not that big of a surprise, considering that you have to live with five people on a few square meters for one and a half week. We ended up with the butjes camper (self-appointed #1), the harem camper, the popcorn camper, and the escalation camper.

During our days in the campers we came across many villages, cities, and beautiful viewpoints. One night we'd stay on top of an uninhabited mountain, while the next we'd stay in the center of a city. One of the fun activities was visiting the Skyline Luge in Rotorua — a Mario Kart style racing track on the side of a mountain.

the prize money didn't last very long


Departure, Groningen


Mount Eden, Auckland


Waiheke Island

We ended up in Wellington after some days of driving. There we visited Victoria University and two companies. Game development studio PikPok had a well-prepared program with short talks by employees from all positions. At crowdfunding platform PledgeMe we drank some very local beer (the brewery was less than hundred meters from the office).

After Wellington we travelled to the South Island. Driving to Queenstown was difficult, since we had to avoid a forbidden dangerous road. We managed to get there and had the honor of enjoying world's best hamburger at Fergburger. After several burgers we moved on.

and the second one was even longer: 18.5 hours

Our journey ended in Christchurch, the city that was destroyed by earthquakes in 2010 and 2011. Christchurch was still a large construction site, which provided a real life experience of the impact of an earthquake. There were few fun things to do, but fortunately we stayed in a fancy hostel where we could enjoy ourselves. We also had our last company visit: we visited Jade Software, New Zealand's biggest software company, and their spin-off Wynyard.

The three weeks passed way too fast. Before we knew it we were back on the plane, en route to the cold Netherlands.

Annet Onnes:
The journey started at a quarter to four for the majority of the group on the 17th of April at the central station of Groningen. I did not join them there however: my journey started just a few hours earlier at the same place. I conveniently lost my debit card earlier that week, so I had to go to my parents' house to pick up my new one. Since my mother is terrified of me ever missing a plane, she decided to drive me to Schiphol so I did not have to depend on the public transport.

I had been on a plane before, but never really for longer than three hours. This time we had to fly twice: one flight of eight hours, and the second one was even longer: 18.5 hours. After getting everyone together, we went to check in our luggage. We all stood in line, being excited about what lay ahead, but some of us also were a bit nervous for the long flights. It turned out I was majorly underpacked, but that was not as bad as being overpacked, like

Rush Digital, Auckland

Motorhomes, near Auckland

Vulcanic activity, Rotorua

Jona. He needed Maikel's help to close his suitcase.

We ate our last non-airplane food for the coming 30 hours or so while waiting to board. We were all seated in one group, although we were unable to figure out how they organized our seats. I got a spot in between Rayan and the aisle. We immediately started playing around with the media systems.

Yes, there are enough films on those things to entertain you for far more than 24 hours, but at some point you really do want to get some sleep. Even though I was exhausted after the first flight to the point of a breakdown on Dubai airport, I only got one hour of sleep.

We were transported with (private!) shuttle busses to the hostel, a bright pink building near the main shopping street of Auckland. All still numb from the flight and not quite realizing that we were actually really in New Zealand, we settled into our rooms and most people got New Zealand sim cards for cheap data. The group split up for dinner, and I ate in the hostel with about half of the group. Afterwards we wandered about the supermarket stunned by the weird, unfamiliar products while trying to figure out the conversion rate of NZD to euros.

After dinner, nobody lasted very long (except for Sanne who has the amazing skill to fall asleep straight after boarding a plane) and we all ended up exhausted in bed around eight o'clock.

Jonathan Maas:
Everyone knows racing is fun. Unfortunately, we don't have that many mountains — let alone racing tracks — in The Netherlands. On the other side of the world, however, the Waiheke Island provides plenty of mountains, roads, and mountain roads to race on. After taking the ferry early in the morning — and nearly losing my ferry ticket to the strong wind — we arrived with beautiful weather and set off to cross the island on mountain bikes. And I can tell you, racing downhill at roughly 60 km/h and taking sharp turns with an everlasting panorama of either mountains, beaches, or seas was massive fun, even though going uphill was less so. Watching the "scooter squad" — the people who hired scooters instead of mountainbikes — going uphill was definitely better than doing so yourself, mainly because Maikel, who would be leading at the base, was surpassed by every scooter when going uphill.

Near the end of the day, I joined up with Arnoud on the way back, and through our combined navigational skills we managed

Mount Ruapehu, Tongariro

Victoria University, Wellington

Ferry Terminal, Wellington

to get helplessly lost (and found ourselves cycling the same 8 km road three times before we found the right direction). There we were: both smartphones pretty much depleted, in the middle of nowhere, opposite side of the world, and cycling on a cliff road in pitch-pitch-black (oh, and with way too many mosquitos) on an island with an estimated total of twenty street lights (and where it was already dark by 6 p.m.). I know there is some god I must praise for Arnoud remembering he had bicycle lights. In the end, we were only two hours late returning the bikes, and even a Fergburger didn't taste as satisfying as the simple whopper I had at the end of that day.

Maikel Grobbe:
Sometimes you have these moments, you are getting older and start forgetting stuff more easily. This isn't always that big of a problem, for example when you forget to put the fuel lid back on or forget to take your medication with you on the airplane. However, forgetting some things will cause serious problems and might even start a war.

One of these things is forgetting to take the Cover flag with you when you leave your hostel room. This did however happen on Friday the 24th of April around 10.00 a.m. and has since led to what we have come to know as 'The Prank War'. Once the flag was forgotten in the hostel room, the 'Escalation' camper took it upon themselves to take the flag and protect it as if it was their own. This led to camper #1, which were the ones who forgot the flag, to launch multiple missions to reclaim their flag. These missions all failed, because of the strict defence system in place (to be more precise, people getting pushy when the attackers got near the hiding place).

a nightly covert operation was launched

After multiple unsuccessful daytime tries, a nightly covert operation was launched. After searching the camper for more than 15 minutes amidst several sleeping defenders, the operation had to be cancelled without result, due to fear of detection. Once camper #1 failed to secure the flag again, drastic measures were taken and the food coloring was introduced. After several members were sent to create a diversion, the rest of the attackers got underway, stole the keys of the 'escalation' camper without them knowing, opened the water hatch and dropped the black coloring in it.

Franz Josef Glacier


Queenstown


Lake Opuha

After this the war had properly started, and the next morning the 'escalation' camper tried to retaliate by throwing a bucket of black water on the windshield. This had little to no effect, causing them to bring out the big guns, or well 'big' 'gun' really, in the form of a ketchup bottle. After they had their joy and their water was cleaned out, a momentary ceasefire was called.

Some nights later, all of a sudden, a weird crackling sound was heard outside of camper #1. Within seconds the defence system was launched and the attacker holding black wrapper foil was comprehended. Camper #1 decided not to further pursue the war, but 'escalation' had one final trick up their sleeves. Once the group had taken residence in a hostel again, a brilliant scheme was taught up to finally utilize the black wrapper foil and cover all of the belongings of a certain somebody. 'Escalation' as of now thinks the war is over and that they have won. Little do they know...

Annet Onnes:

Our time in New Zealand seemed to have passed within the blink of an eye when it was suddenly time to pack our suitcases again and hop on the bus to the airport of Christchurch. We were now all familiar enough with long flights, so we discussed which films we were going to watch and during what part of the flight it would be best to try and sleep. It turned out this was between Bangkok and Dubai, when the lights were turned down and we could see small lights like a starry night on the roof of the plane. I spent my time catching up on the reading work for a course, like some others who were studying for Biopsychology. Many films were watched, like Ex Machina, a film you should definitely watch as an AI/CS student.

Martijn got all together with "milliemilliemillie"

After 5 chapters, 3 films and countless games of patience, we landed in Amsterdam, finally back in our little flat country. Quickly switching our sim cards to tell our family, friends and significant others that we got there safely. We picked up our luggage and then some of us finally got their dramatic 'hello goodbye' moments. As a final goodbye to our time in New Zealand, Martijn got all together with "milliemilliemillie" and those who were there know what followed...

Jade Software, Christchurch

Cardboard Catherdral, Christchurch

Surfing class, Christchurch










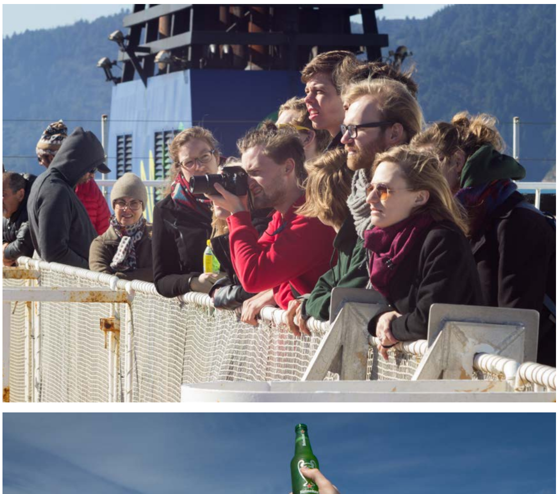



STUDIEVERENIGING
Cover
apollo


43KM
OF PURE


KAWARAU & SHOTOVER RIVERS

COLOPHON

The Brainstorm is a magazine published by study association Cover and is distributed among its members, staff members and other interested people. The Brainstorm comes out at least three times a year in an edition of 500.

Editors

Chairman	Ben Wolf
Secretary	Annet Onnes
Treasurer	Arnoud van de Meulen
Senior Editor	Steven Warmelink
Junior Editor	Sanne Bouwmeester
Layout	Ben Wolf Annet Onnes

Contact

Study association Cover
attn. The Brainstorm
PO box 407
9700 AK Groningen
brainstorm@svcover.nl
www.svcover.nl