

ACOUSTIC MONITORING OF BATS WITH SELF-ORGANIZING MAPS

word count: 27693

Arne Deloose

Student ID: 0130272

Supervisor(s): prof. dr. Bernard De Baets and dr. ir. Michiel Stock

Tutor(s): ir. Laura Hoebeke

A dissertation submitted to Ghent University in partial fulfilment of the requirements for the degree of master in Bioscience Engineering.

Academic year: 2018 - 2019

De auteur en promotors geven de toelating deze scriptie voor consultatie beschikbaar te stellen en delen ervan te kopiëren voor persoonlijk gebruik. Elk ander gebruik valt onder de beperkingen van het auteursrecht, in het bijzonder met betrekking tot de verplichting uitdrukkelijk de bron te vermelden bij het aanhalen van resultaten uit deze scriptie.

The author and promoters give the permission to use this thesis for consultation and to copy parts of it for personal use. Every other use is subject to the copyright laws, more specifically the source must be extensively specified when using results from this thesis.

Gent, June 7th, 2019

The promoters,

prof. dr. Bernard De Baets

dr. ir. Michiel Stock

The author,

Arne Deloose

DANKWOORD

Deze thesis had niet tot stand kunnen komen zonder de medewerking van een hele groep mensen, die ik dan ook graag bedank in dit deel. Allereerst mijn promotors Michiel Stock en Bernard De Baets. Zij hebben het onderwerp aangeboden en de samenwerking met INBO mogelijk gemaakt. Ook voor advies en begeleiding waren zij steeds beschikbaar, alsook het doorsturen van de vele interessante referenties over vleermuizen en dataverwerking. Ook wil ik mijn tutor, Laura Hoebeke, bedanken om elke week tijd vrij te maken voor mijn vele vragen en voor het lezen en verbeteren van de vele versies van deze dissertatie.

Verder wil ik het INBO bedanken voor het beschikbaar stellen van het onderwerp en de bijhorende data, alsook voor de uitleg rond vleermuizen en geluidsverwerking. Veel van deze data werd in hun vrije tijd verzameld en daarom verdienen deze mensen zeker een extra bedanking. In het bijzonder wil ik Thierry Onkelinx bedanken, die mijn thesis begeleidde vanuit het INBO.

De mensen van Natuurpunt, voor het organiseren van de 'Nacht Van De Vleermuis', wat een zeer leerrijke avond was waar ik veel heb geleerd. Ook voor de vele informatie over vleermuizen waarvan ik nuttig gebruik heb gemaakt, en de aanbeveling van het zeer nuttige boek 'Veldgids-Vleermuizen van Europa' van Dietz and Kiefer.

Deze thesis maakte gebruik van vrij veel software. Daarom zou ik graag de onderzoekseenheid KERMIT, met in het bijzonder Joris Meys en Jan Roels, bedanken voor hun uitleg rond Python, LaTeX and GitHub. Verder nog Bac Nguyen Cong voor de code rond distance metric learning.

Als laatste, Niels Hilderson en mijn familie voor de vele morele steun tijdens het schrijven en programmeren.

Arne Deloose

Juni 2019

CONTENTS

Dankwoord	i
Contents	vi
Nederlandstalige samenvatting	vii
Summary	ix
Graphical abstract	xi
List of abbreviations	xiii
1 Introduction	1
1.1 Problem statement	1
1.2 Objectives	1
1.3 Outline	2
2 Ecology and traits of bats	3
2.1 What are bats?	3
2.2 Species and evolution	4
2.2.1 Species of bats	4
2.2.2 Bats in Belgium	6
2.3 Ecological and economical importance	8
2.4 Echolocation	10
2.4.1 Overlap with other senses	10
2.4.2 Mechanism	11
2.4.3 Limitations of echolocation	12
2.4.4 Evolution of echolocation	13
2.4.5 Classification of echolocation	14

3	Conservation and monitoring of bats	17
3.1	Threats	17
3.2	Conservation measures	20
3.3	Legislation	20
3.3.1	European Union	20
3.3.2	Belgium	21
3.4	Non-governmental organizations	22
3.5	Current status	23
3.5.1	Worldwide	23
3.5.2	European Union	24
3.5.3	Belgium	24
3.6	Monitoring	25
3.6.1	Non-audio methods	25
3.6.2	Audio methods	26
4	Ecoacoustics	27
4.1	What is ecoacoustics?	27
4.2	Acoustic niche hypothesis and acoustic adaptation hypothesis	27
4.3	Applications of ecoacoustics	28
4.4	Bat detectors	29
4.4.1	Heterodyne detectors	30
4.4.2	Frequency division detectors	31
4.4.3	Digital detectors	31
4.5	Commercial software for classification of bats	33
4.6	Technical problems with bat detectors	34
4.6.1	Recording	34
4.6.2	Variation in detectors	35
4.6.3	Anomalies	35
4.7	Challenges with bat sounds	36
4.7.1	Convergence	36
4.7.2	Variation in calls	36
4.7.3	Interference	37

5	Machine learning for data analysis, classification and clustering	39
5.1	What is machine learning?	39
5.2	Types of machine learning	40
5.2.1	Supervised learning	40
5.2.2	Unsupervised learning	40
5.2.3	Semi-supervised learning	41
5.3	Clustering methods	41
5.3.1	<i>K</i> -means clustering	41
5.3.2	Self-organizing maps	42
5.4	Distance metric learning	44
5.5	Data transformation	46
5.5.1	Spectrograms and Fourier transformation	46
5.5.2	Scaleograms: wavelet transformation	47
5.6	Unbalanced data	48
6	Pre-processing and clustering	51
6.1	Data description	51
6.2	Pre-processing	53
6.2.1	Transformation	53
6.2.2	Region of interest extraction	54
6.2.3	Feature extraction	54
6.2.4	Reference library	58
6.3	Analysis	58
6.3.1	Exploration methods	59
6.3.2	Classification method: self-organizing map	59
6.3.3	Visualization tool	61
7	Results and discussion	63
7.1	Application of the tool	63
7.2	Evaluation experiments	66
7.2.1	Description of the datasets	66
7.2.2	Method: <i>K</i> -nearest neighbors and Kappa statistic	66
7.2.3	Self-organizing maps	67
7.3	Experiments	68

7.3.1 Experiment one: generalization within species	68
7.3.2 Experiment two: influence of K and DML	69
7.3.3 Experiment three: generalization to different species	71
8 Conclusions and future perspectives	73
Bibliography	75

SAMENVATTING

In Vlaanderen zijn er 18 soorten vleermuizen. Al deze soorten zijn beschermd onder Europese en lokale wetgeving. Onder deze wetgeving moeten de populaties elke zes jaar gerapporteerd worden. Dit is vrij moeilijk gezien vleermuizen nachtdieren zijn en dus moeilijk visueel te monitoren zijn.

Vleermuizen kunnen vliegen en jagen in complete duisternis. Hiervoor gebruiken ze echolocatie: ze zenden geluid uit en luisteren zorgvuldig naar de echo's die terugkomen. Met deze techniek kunnen vleermuizen hun omgeving in kaart brengen. Verschillende vleermuissoorten gebruiken verschillende geluiden, wat het mogelijk maakt om deze geluiden als een monitoring techniek te gebruiken. Hiervoor worden hoogfrequente opnames gemaakt en de resulterende spectrogrammen worden door experts geanalyseerd. Eén enkele nacht op één locatie kan tot 1500 opnames bevatten, wat manuele labeling moeilijk maakt.

Data van het Instituut voor Natuur en Bos Onderzoek (INBO) wordt gebruikt om een praktische tool te ontwikkelen. Deze tool maakt gebruik van self-organizing maps. Deze maps maken gebruik van neuronen waarbij elk neuron een groep datapunten beschrijft. Het resultaat hiervan is dat elke vleermuissoort gelinkt kan worden aan een of meerdere neuronen. Als de geluiden reeds geclassificeerd waren, kan een neuron gelinkt worden met één of meerdere vleermuissoorten. Als dit niet het geval is, kunnen enkele geluiden manueel gelabeld worden en op basis hiervan kan een conclusie getrokken worden voor alle geluiden die matchen met ditzelfde neuron. De data wordt dus verdeeld in groepen. De analyse kan dan op groepsniveau gebeuren in plaats van op het niveau van individuele geluiden, wat de analysetijd sterk verkort. Het doel van deze tool is dus niet de huidige analyse te vervangen, maar eerder om deze efficiënter te laten verlopen. Hiermee kan een betere monitoring van vleermuispopulaties gebeuren.

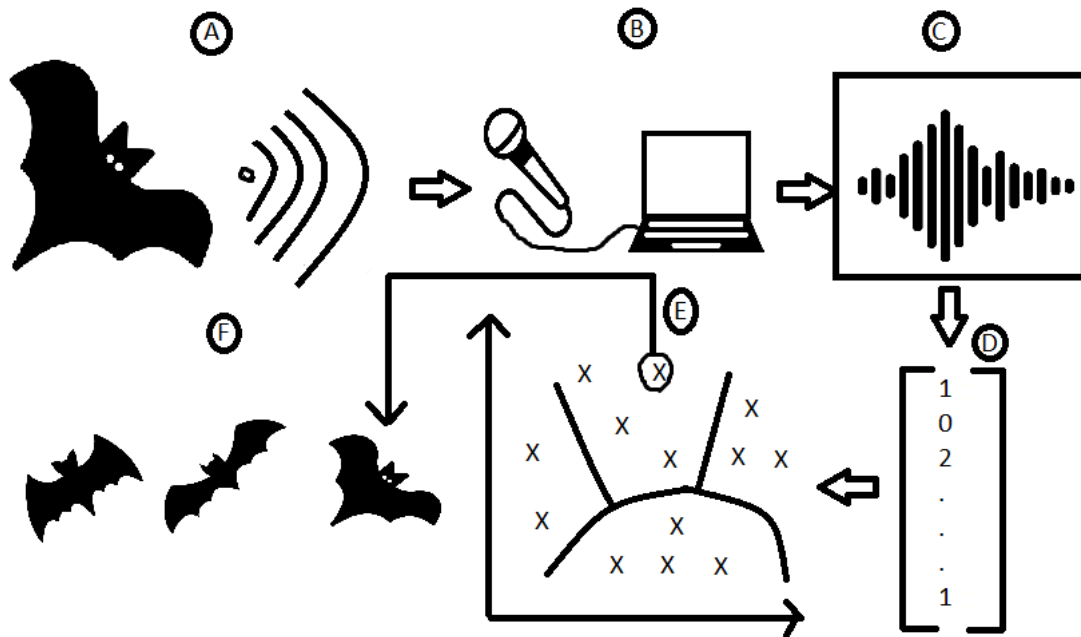
SUMMARY

In Flanders there are 18 species of bats. All of these species are protected under European and local laws. Under these laws, the populations must be reported every six years. Due to the nocturnal nature of bats, visual monitoring is rather challenging.

Bats can fly and hunt in total darkness. To that end, they use echolocation: they emit sounds and listen carefully to the echoes that return. Using this technique, bats can map their surroundings. Different bat species use different sounds, which makes it possible to use these sounds as a monitoring technique. To that end, high-frequency recordings are made and the resulting spectrograms are analyzed by experts. For a single night in a single location, this can result in more than 1500 recordings, which makes full manual labeling difficult.

Data from the Research Institute for Nature and Forest (INBO) is used to develop a practical tool. This tool uses self-organizing maps. These maps use neurons that describe a group of datapoints. The result of this tool is that each bat species is tied to one or more neurons. If sounds were already classified, a neuron can be tied to one or more bat species. If not, a few sounds can be classified manually and from this a conclusion can be drawn for all the sounds matching with this neuron. In other words, the data is divided into groups. The analysis can then take place at the level of these groups instead of the level of individual sounds, which shortens the time needed dramatically. The goal of this tool is not to replace the current analysis, but to make it more efficient. This can lead to a better monitoring of bat populations.

GRAPHICAL ABSTRACT



There are six steps in the analysis of bat sounds. A certain bat emits a sound (A). This sound is then recorded by a bat detector (B). Based upon this sound, a spectrogram is created and the bat sound is isolated (C). From this spectrogram, features are extracted (D). These features are numbers that describe certain aspects of the spectrogram. Then, this datapoint is plotted on a self-organizing map (E). Based upon how similar the datapoint is to other datapoints, it takes a specific place in the map. Datapoints closer to it are more similar while those farther away are less similar. An expert then looks at the different regions in the map and ties them to one or more species (F). (Image source: Noun Project: rivercon (first bat), andriwidodo (microphone), worker (laptop), maxim kulikov (soundwave), leona grande (middle bat) and tulpahn (left bat))

LIST OF ABBREVIATIONS

AAH: acoustic adaptation hypothesis

ANH: acoustic niche hypothesis

ATL: Atlantic region

BMU: best matching unit (SOM)

CE: critically endangered (Red List IUCN)

CON: continental region

DD: data deficient (Red List IUCN)

DML: distance metric learning

DMLMJ: distance metric learning through Jeffrey divergence

E: extinct (Red List IUCN)

EEA: European environment agency

EN: endangered (Red List IUCN)

EU: European Union

EW: extinct in the wild (Red List IUCN)

FD: frequency division bat detector

INBO: research institute for nature and forest

IUCN: international union for conservation of nature

KNN: *K*-nearest neighbors

LC: least concern (Red List IUCN)

MDS: multi-dimensional scaling

MSE: mean squared error

NE: not evaluated (Red List IUCN)

NT: near threatened (Red List IUCN)

ROI: region of interest

SOM: self-organizing map

SSIM: structural similarity index

TDSC: time domain signal coding bat detector

TE: time expansion bat detector

t-SNE: t-distribution stochastic neighboring embedding

VU: vulnerable (Red List IUCN)

ZCA: zero-crossing analyzer (bat detector)

CHAPTER 1

INTRODUCTION

Acoustic bat detectors are one of the most popular monitoring methods used in the study of bats, mainly to determine their presence, species and activity. The main advantage over other methods such as telemetry, light traps and winter counting, is their non-invasive nature and low labor cost. Bat detectors can be left out in the field to passively collect data for days or even weeks, which is generally not possible with other methods.

1.1 Problem statement

The main bottleneck for bat detectors is the annotation of recordings and the large amounts of data collected. A single night can lead to hundreds of audio files that need to be annotated manually. Because the annotation requires expert knowledge, citizen science is generally not an option. The Bat Detective project¹ uses citizen science, but the main goal of this project is to separate bat calls from non-bat calls, not to classify individual bats. Annotation can also be done using software such as SonoBat or BatSound, but the accuracy of these packages is limited, they can be quite expensive and they are generally not adapted to a specific region. Currently, Belgium does not have a practical tool adapted to the region. All reporting happens with manual annotation. Legislation at the level of the European Union requires member states to estimate populations of bats and report these every six years.

1.2 Objectives

The objective of this research was to build a tool that can help with the annotation of bat sounds. There were not enough labeled data available to build an automated classifier. However, there were partially labeled data provided by the Research Institute for Nature and Forest (INBO). Based on these data, a tool was built to assist the manual classification.

¹<https://www.batdetective.org/>

1.3 Outline

In Chapters 2 and 3, bats are discussed with a focus on the species, ecology, legislation, threats and conservation. Furthermore, monitoring methods and echolocation are discussed. This is important to provide the proper context of the problem. Chapters 2 and 3 make clear why bats need to be monitored, the basis of the monitoring techniques and the challenges associated with them. Chapter 4 focuses on ecoacoustics, the use of audio data for monitoring purposes. Ecoacoustics is discussed in general and with a focus on bat applications. This provides the proper context surrounding audio data and highlights the potential challenges down the road. Chapter 5 discusses machine learning, a discipline of computer science that can assist the analysis of the audio data from bats. These techniques are then applied to a dataset provided by INBO (Research Institute for Nature and Forest). This is discussed in Chapters 6 and 7. Finally, Chapter 8 contains the conclusions and future perspectives.

All code used is available on Github².

²<https://github.com/ArneDeloose/Masterproef2018Code>

CHAPTER 2

ECOLOGY AND TRAITS OF BATS

This chapter will focus on four subjects. Firstly: what are bats and what are their most important traits? Secondly: what are the different phylogenetic groups and species of bats in the world and in Belgium? Thirdly: what is the ecological importance of bats? And lastly, what is echolocation? Echolocation is one of the most important traits of bats and this will play a vital role in the way they are monitored (see Chapter 4).

2.1 What are bats?

What are bats exactly? Formally speaking, bats are all mammals that belong to the order of the Chiroptera. These animals comprise about 20% of all classified mammal species, making them the second largest order, right after the Rodentia (Wilson, 2005). Bats have three important traits that distinguish them from other mammals.

The first trait is also their namesake. 'Chiro' comes from the Greek 'cheir' which means hand and 'pteron' means wing (Merriam-Webster, 2019). In other words: a bat is a mammal that has transformed its hands into wings by using a membrane (*patagium*). This is not unique to bats. Several other species such as flying squirrels have similar membranes. However, bats are the only mammals that can use their wings for powered flight rather than simple gliding. Gliding means that an external force is needed to ascend. Powered flight means the wings provide the necessary lift on their own. Only three other groups of animals evolved flight: insects, birds and pterosaurs. Pterosaurs are sometimes excluded because they have been extinct for millions of years and their flight mechanics are not well understood (Dietz and Kiefer, 2017; Altringham, 2011; Sato et al., 2009).

A second important trait of bats is their nocturnal nature. Most bats are nocturnal animals flying in twilight and darkness. Most likely, bats are nocturnal to avoid competition for food from other animals. Other theories are to avoid predators and to avoid overheating during the day (Speakman, 2001; Mikula et al., 2016; Voigt and Lewanzik, 2011). During the day, bats roost. Characteristic of bats is that they always hang upside down when roosting. This allows them to take flight very fast by letting themselves fall down (Dietz and Kiefer, 2017).

The last important feature is echolocation, the ability to locate and identify objects using sound. Bats are most notable in using this technique, although not all bats use it and several other species such as dolphins, shrews and birds use a similar technique (Holland et al., 2004).

2.2 Species and evolution

2.2.1 Species of bats

Bats come in various sizes. At the highest end, there is the giant golden-crowned flying fox, which can reach wingspans of 1.7 m and weigh up to 1.6 kg (Nowak, 1999). At the lowest end there is Kitti's hog-nosed bat, which has wingspans of only 15 cm and weighs only 2 g (Nowak, 1999). This size difference comes with a wide range of other differences. Because of these differences, bats are traditionally divided into two groups: megabats and microbats (Prothero, 2017).

As their name suggests, megabats are usually larger than microbats, but there are other differences as well. Some differences are morphological, such as the absence of a tail, tragus and noseleafs and the presence of a claw on the second finger. A tragus and noseleafs both assist in echolocation, which is discussed in Section 2.4.2. Megabats do not need a tragus, because they usually do not have echolocation. Notable exceptions are the ten species of the *Rousettus* genus that can use tongue clicking noises as a primitive form of echolocation. Furthermore, research by Boonman et al. (2014) suggests there are megabats that use the sound of their own wings for echolocation. However, true echolocation is most likely not possible due to their larger body size. This is discussed further in Section 2.4. Instead, megabats rely on different senses, mainly smell and sight. This shift in senses comes with relatively larger eyes compared to microbats. The use of their senses also matches their diet. The usual diet of megabats consists of fruits and nectar. Since these are stationary, an advanced system like echolocation is not necessary. Megabats are a small group. Only 187 species of megabats are listed on the IUCN Red List (International Union for Conservation of Nature), compared to 1244 bat species in total. The IUCN Red List is the standard reference work to assess the current conservation status of all species. Conservation is discussed in detail in Section 3.5. All megabats belong to one family: the Pteropodidae. Common synonyms for megabats are flying foxes and Old World fruit bats. The first name refers to their typical faces and size and the second name refers to their habitats. Megabats are only found in the tropical and subtropical regions of Eurasia, Oceania and Africa. Since our research focuses on Belgium, they will not be discussed any further. For more information on megabats, we refer to Neuweiler (2000).

Microbats are much smaller. Their diet usually consists of insects, with some bigger bats hunting frogs, lizards, fish or even other bats. A notable exception are the three vampire bat species that feed on animal blood instead. Microbats are a larger group with around 1000 species (IUCN). The biggest groups within the microbats are the vespers or evening bats with around 400 species and the Phyllostomidae or New World leaf-nosed bats with around 200 species (IUCN). Microbats always have echolocation. They also have relatively larger ears and smaller eyes than megabats. Most of them have a tail, which is always absent in megabats. Microbats are found all over the world except for the coldest regions. While they are found everywhere, biodiversity is centered around tropical and subtropical areas (Dietz and Kiefer, 2017). An illustration of this is given in Figure 2.1.

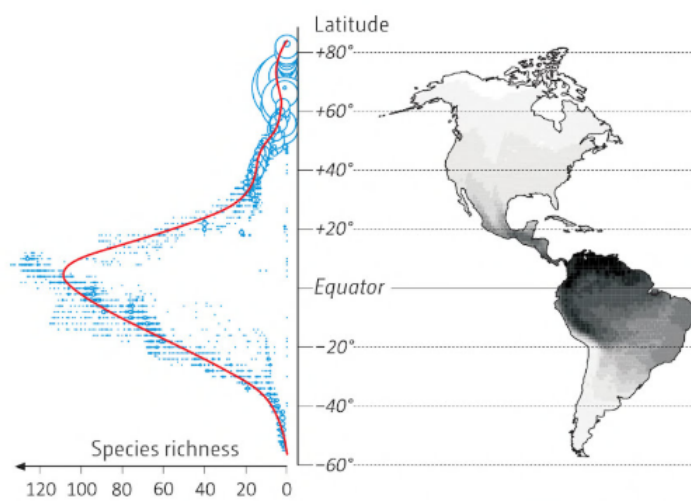


Figure 2.1 Distribution of bat species per latitude in America. Cells of 100 km x 100 km are used. The size of the circles indicates the number of cells with the same number of species. Tropical areas have the highest biodiversity. Source: Dietz and Kiefer (2017).

The classification in megabats and microbats is based on morphology, diet and behavior. However, modern classification is increasingly based on genetic data and evolution instead. New research suggests that there are five microbat families that are genetically much closer to megabats. Therefore, a new classification was proposed: Yinpterochiroptera and Yangochiroptera. Yangochiroptera contains the remaining microbats. Yinpterochiroptera splits immediately into Pteropodidae and Rhinolophidae. The first group contains the original megabats, the second group contains five families that were originally microbats (Springer et al., 2001). An illustration of this can be found in Figure 2.2.

All of this is part of a wider debate on the monophyletic nature of bats. A monophyletic group is a group with a single common ancestor. Polyphyletic groups have multiple ancestors. Currently, it is not yet clear whether bats are monophyletic. It is entirely possible that megabats evolved from a different ancestor than microbats. The current theories suggest that flight developed only once and the two groups split about 60 million years ago. Cer-




Figure 2.2 New classification of bats.

Frequencies from: Obrist Martin K. (2007), the value of the Pond bat was missing.

Name (English)	Name (Scientific)	Frequency call (kHz)	IUCN rating
Bechstein's bat	<i>Myotis bechsteinii</i>	27 – 104	NT
Brandt's bat	<i>Myotis brandtii</i>	28 – 104	LC
Pond bat	<i>Myotis dasycneme</i>	–	NT
Daubenton's bat	<i>Myotis daubentonii</i>	27 – 81	LC
Geoffroy's bat	<i>Myotis emarginatus</i>	36 – 113	LC
Greater mouse-eared bat	<i>Myotis myotis</i>	22 – 86	LC
Whiskered bat	<i>Myotis mystacinus</i>	28 – 100	LC
Natterer's bat	<i>Myotis nattereri</i>	14 – 109	LC
Barbastelle	<i>Barbastella barbastellus</i>	26 – 48	NT
Serotine bat	<i>Eptesicus serotinus</i>	22 – 47	LC
Lesser noctule	<i>Nyctalus leisleri</i>	22 – 49	LC
Common noctule	<i>Nyctalus noctula</i>	18 – 34	LC
Nathusius' Pipistrelle	<i>Pipistrellus nathusii</i>	36 – 62	LC
Common pipistrelle	<i>Pipistrellus pipistrellus</i>	43 – 74	LC
Brown long-eared bat	<i>Plecotus auritus</i>	23 – 56	LC
Grey long-eared bat	<i>Plecotus austriacus</i>	18 – 45	LC
Greater horseshoe bat	<i>Rhinolophus ferrumequinum</i>	69 – 84	LC
Lesser horseshoe bat	<i>Rhinolophus hipposideros</i>	90 – 111	LC

Table 2.1 Overview of the bat species present in Belgium.

IUCN ratings: LC: least concern, NT: near threatened

tain megabats later developed traits of microbats and that is why they were misclassified (Springer et al., 2001).

This classification issue remains an unsolved problem. The problem relates to fossil data. All flying animals have brittle skeletons to reduce weight. As such, bones do not fossilize well. Research by Eiting and Gunnell (2009) estimates that only 12% of bat genera have a fossil record. This makes accurate classification rather difficult.

2.2.2 Bats in Belgium

The primary focus will be on the bats of Belgium. An overview of species is given in Table 2.1 and the guilds are discussed further in Section 2.4.5. There are a total of eighteen species in Belgium. This number is based on the species that are reported to the European Union (EU). However, several sources claim that there are in fact more species (Wilson (2005), Dietz and Kiefer (2017), Natuurpunt and IUCN data). The largest estimate would be twenty-four species. We will first discuss the formal eighteen species and then move on to borderline cases. Of the eighteen species, sixteen belong to the vesper bats. Vespers are mostly insectivores, with some exceptions eating small birds or fish. They have relatively large ears and small noses.

Within Belgium, the *Myotis* genus has the most representatives with a total of eight species. *Myotis* bats are commonly called mouse-eared bats because they have long, narrow ears. They usually fly in open spaces and have a high-range echolocation. An example of a *Myotis* bat is the whiskered bat, shown in Figure 2.3a.

Continuing, there is one species of the *Barbastella* genus: *Barbastella barbastellus*. This species is shown in Figure 2.3b. It is a rare species, so rare it was considered extinct in Belgium until 2014, when Natuurpunt reported new sightings. They roost in damaged trees, which become increasingly rare due to forest management. *Barbastella* species have small eyes, a small nose and large, broad ears.

The next genus is the *Eptesicus* genus with only the serotine bat, shown in Figure 2.3c. While this is an abundant species, populations are declining in many areas due to construction works and better isolation. *Eptesicus* bats have buildings as their habitat. The serotine bat has black ears and nose and is quite large, with wingspans of up to 38 cm. It shows morphological similarities with the common noctule, but distinguishes itself rather easily by flying out after sundown, while the common noctule flies out before sundown.

The next species are the common noctule and lesser noctule of the *Nyctalus* genus. The common noctule is shown in Figure 2.3d. As was already mentioned, morphologically speaking, these bats are similar to the *Eptesicus* genus, but they show different behavior. Apart from the difference in emergence, they also have different winter tactics. Rather exceptionally, the common noctule shows female bias migration. The females migrate south in winter, while the males hibernate.

Next, there is the *Pipistrellus* genus with two species. The common pipistrelle is the most abundant species of Belgium (according to Natuurpunt). *Pipistrellus* species are commonly urban species. They hunt in parks or gardens and live in buildings. A picture of a common pipistrelle is shown in Figure 2.3e.

Lastly, there are the *Plecotus* or long-eared bats. Two species of this genus are present in Belgium. As the name suggests, these bats have very long ears. They use relatively low frequencies for echolocation, which have a long range. These species are not threatened now, but there are concerns that climate change could affect that. They are known for their long lifespans, low reproduction rates and small population, so their ability to adapt to rapid changes is low. An example of the *Plecotus* genus is the *Plecotus auritus*, depicted in Figure 2.3f.

This leaves only two species that are not vespers. These last two are the *Rhinolophus ferrumequinum* and *Rhinolophus hipposideros* from the Rhinolophidae family, commonly called the horseshoe-bats. This is one of the families that was added to the Yinpterochiroptera, which means that these two species probably evolved from megabats. Character-

istic of these species are their horseshoe-shaped noseleaves that help with echolocation. As we can see in Table 2.1, the two *Rhinopholus* species have a much higher echolocation frequency. The *Rhinolophus ferrumequinum* is depicted in Figure 2.3g.

Apart from these bats, there are other bats that have an unclear status. It is not always straightforward to determine whether a bat is present in a certain country. As an example, the Alcathoe bat will be discussed (*Myotis alcathoe*). This species is morphologically almost indistinguishable from the whiskered bat, but was recognized as a separate species based on genetic data in 2001. Confirmed recordings show a patchy distribution with many isolated records (Niermann et al., 2007). As such the IUCN lists its status as data deficient. In Belgium, there are only thirteen recordings in the Walloon region (Nyssen et al., 2015). All of this makes it difficult to conclude unambiguously whether the Alcathoe bat is permanently present in Belgium or not. Apart from the Alcathoe bat, there are three other rare bats in a similar situation. The greater noctule (*Nyctalus lasiopterus*), the parti-coloured bat (*Vespertillio murinus*, shown in Figure 2.3h) and Kuhl's pipistrelle (*Pipistrellus kuhlii*) are all examples of rare bats with only a handful of isolated recordings.

This leaves two more bats: the soprano pipistrelle and the northern bat. Both of these species are widespread in Eurasia, but do not have colonies in Belgium. However these bats are present in parts of Germany and France that are very close to the border. Bats can cover large distances in a single night. This makes it possible that while these bats do not have permanent colonies in Belgium, they occasionally pass through and are detected this way.

2.3 Ecological and economical importance

Bats have several important ecological and economic functions. For a full overview we refer to Kasso and Balakrishnan (2013). The four most important ones will be discussed here.

In tropical climates, megabats play a vital role in flower pollination and seed dispersal. With motion cameras, the main pollinator of the Australian baobab was shown to be the black flying fox (Groffen et al., 2016). According to Hodgkison et al. (2006), 13.7% of trees in Malaysia relies at least partially on bats for pollination and seed dispersal.

Microbats on the other hand are important for the control of insect populations. In research by Boyles et al. (2011), it is estimated that the loss of bats would cost agriculture 3.7 to 53 billion dollars per year in extra pesticides and yield losses in the United States alone.

The excrement of bats can be collected as a fertilizer, called guano. Due to their short digestion process, faeces from bats are rich in nitrates, phosphates and salts. Today, this is still used in organic farming. In the past, it was also used to make gunpowder as a



(a) *Myotis mystacinus* (Whiskered bat).



(b) *Barbastella barbastellus* (Barbastelle).



(c) *Eptesicus serotinus* (Serotine bat).



(d) *Nyctalus noctula* (Common noctule).



(e) *Pipistrellus pipistrellus* (Common pipistrelle).



(f) *Plecotus auritus* (Brown long-eared bat).



(g) *Rhinopholus ferrumequinum* (Greater horse-shoe bat).



(h) *Vespertillio murinus* (Parti-colored bat).

Figure 2.3 Bat genera in Belgium. All eight genera are represented by one species from that genus. Source: Natuurpunt (Hugo Willocx and Yves Adams).

source of saltpeter. Around 1850, the US imported 760 000 tonnes of guano per year from various islands in the tropics (Smil, 2004). In 1856, the US wrote the Guano Islands Act, a federal law that allowed any US citizen to claim islands with guano (as long as they were not occupied or claimed by a different country). It even empowered the president to use the military to protect these islands. Roughly a hundred islands were claimed this way in the past, with ten islands still being part of the US today (although several claims are being disputed by Colombia and Haiti). Guano also plays a role in soil fertility and nutrient transfer within an ecosystem.

Lastly, bats can be used as a bioindicator. Bioindicators are organisms whose population or activity is indicative of the state of the environment. The pollution state of a river for example can be assessed through the absence or presence of certain invertebrates. Jones et al. (2009) report several stress factors that can lead to a change in population or activity of bats. These stress factors are: extremes of drought, heat, cold and precipitation, cyclones and sea level rise, deterioration of water quality, agricultural intensification, loss and fragmentation of forests and pesticide use. This way, the presence of bats can be used as a proxy for other problems, mainly climate change. de Oliveira et al. (2017) used this technique successfully as a proxy for habitat destruction through deforestation in Brazil.

For countries in temperate climates such as Belgium, insect control is the most important service that bats provide. The other three ecological functions are more important in tropical climates (Kasso and Balakrishnan, 2013).

2.4 Echolocation

The word sonar is short for sound navigation ranging, the use of sound to detect objects (Encyclopaedia Britannica 2019). It is a technique that is mainly used in submarines. When an animal uses this technique, it is called bio-sonar or animal echolocation. Microbats are one of the most notable groups that use this technique. Most megabats do not use echolocation at all. A notable exception is the genus *Rousettus* which uses a primitive form of echolocation (Holland et al., 2004).

2.4.1 Overlap with other senses

Echolocation is the main sense microbats rely on for hunting, communicating and navigation. However, it is not the only sense they have. Despite popular belief, all bats have vision. This vision is usually mesopic (light/dark), but some bats can see in UV as well (Müller et al., 2009). Bats also have magnetoreception. Contrary to birds, bats have polarity-based reception. Birds estimate their latitude based on the strength of the magnetic field, bats

differentiate north and south. Research by Tian et al. (2015) shows that bats can use their magnetoreception even at only one fifth of the normal magnetic field strength. These other senses can provide extra information that is combined with information from echolocation.

Apart from the overlap between senses, there is also a physical overlap. The muscles used in flight, breathing and echolocation can overlap. Because of this, echolocation is much more common in flying animals (Speakman et al., 1989). Altringham (2011) suggests that this coupling is more difficult for larger animals. This would make echolocation less interesting for larger bats in terms of energy. However, research by Speakman and Racey (1991) has found that the energy costs for flight are so high that echolocation is not significant in comparison. It is possible that the main reason why megabats do not have echolocation is simply because their diet does not require them to.

2.4.2 Mechanism

Echolocation works by emitting sounds from the larynx (voicebox) and analyzing the returning echoes. From these echoes, three things can be determined. The first metric is distance to the object. This is determined from the time delay between pulse and echo. The speed of sound is around 340 metres per second. If the time between a call and the echo is one second, the object is around 170 metres away. This 170 metres is then corrected for the speed of the bat itself. If a bat is flying towards the object, the echo will return slightly faster. The faster they fly, the bigger this effect becomes (Jones, 2009). The elevation of the object can also be estimated. This is determined from interference patterns on the tragus. The tragus is a circular bump that partly covers the ear canal (Figure 2.4). Horseshoe bats are able to move their ears vertically independent of one another. This results in different intensities at each ear which can be used to calculate elevation (Müller, 2004; Jones, 2009). Lastly, the direction of the sound can be estimated. This is determined from the intensity difference between both ears. When the object is to the left, the intensity at the left ear will be higher (Jones, 2009).




Figure 2.4 Tragus of a bat ear. Source: Chiu and Moss (2005).

The intensity of the sounds can be a problem. Bat calls can have intensities of up to 140 decibels, louder than a military jet aircraft (Surlykke and Kalko, 2008). This means that in theory, bats would damage their own hearing. However, they have developed two systems to counter this. Firstly, bats are capable of contracting the muscles in their middle ear which makes them temporarily deaf (Teeling, 2009). In low-duty echolocation, the time delay in call and echo is large enough to allow for contraction and relaxation of these muscles. However, in high-duty mode, bats emit continuous calls which make this impossible. High-duty mode is mostly used when chasing an insect. This mode requires even more advanced systems. In high-duty mode, bats use the Doppler effect (Jones and Holderied, 2007). They know exactly how fast they are flying and adapt the frequency of their calls to this flight speed. However, their ears are tuned to a different frequency range. Due to the Doppler effect, returning echoes will have a higher frequency.

Apart from the echoes of their own sounds, bats can also detect the noises produced by certain insects. Certain moth species take advantage of this. They have a tympanum (hearing organ) that detects bat calls. When they hear calls, they produce sounds designed to interfere with echolocation (Hoy and Fay, 1998). Other common responses for insects are twitching of wings to produce random evasive movements. Butterflies are known to use this technique. For a full overview of insect defensive mechanisms to bats, we refer to Miller and Surlykke (2001).

The frequency range of microbat calls in Belgium is between 14 kHz and 113 kHz (see Table 2.1). It is possible bats go even higher, but as we will see later (Section 4.4), most detection devices do not measure above 120 kHz. The range of human hearing is 20 Hz to 20 kHz, so most calls cannot be heard by humans. This is referred to as ultrasound (Cutnell and Johnson, 1997).

2.4.3 Limitations of echolocation

There are several limitations to echolocation. The most important ones will be summarized here. A full overview is given in Schnitzler and Kalko (2001).

Firstly, there are clutter echoes. These are noise signals that come from the surroundings. A classic example would be a butterfly flying between the leaves of a tree. Those leaves will also create echoes. If those echoes interfere with the echoes of the butterfly, it is called backwards masking (Schnitzler and Kalko, 2001).

Secondly, there is forward masking. Here, the cause is internal. The brain of bats needs to do two things: send out calls and interpret the echoes. As long as these two activities are separated well enough in time, there is no problem. But if the prey gets closer and closer, these activities can start to overlap. This is similar to how it is difficult for humans

to talk and listen at the same time. It is also possible that the call overlaps with the echo creating interference. Both make analysis difficult. We call this forward masking, because the interference happened before the echo returned. In backwards masking, the interference happens after the echo has returned. According to Kober and Schnitzler (1990), the minimum detection distance is about 17 cm for every ms of signal. If the distance gets smaller, forward masking is too strong to properly detect an object (Schnitzler and Kalko, 2001).

Thirdly, echolocation has a limited range. As sound waves travel through air, they get absorbed. At some point, they will be too quiet to properly detect. Kober and Schnitzler (1990) calculated that the maximum range of echolocation is on average about 10 m. This range is dependent on many factors such as signal frequency, prey size, humidity and temperature.

2.4.4 Evolution of echolocation

Echolocation plays a vital role in hunting preys. However, it also assists in flight and communication between bats (Schnitzler et al., 2003). Because of this, there are different theories on the initial function of echolocation and its evolution.

The first theory is the flight first theory. Under this theory, flight developed only once in bats. A primitive system of echolocation evolved later to help with orientation. Bats who relied solely on sight would have trouble hunting at night, bats who use echolocation can hunt whenever they want to. At night, there is less competition for food, hence bats with echolocation could take this niche easily. Another possible reason to become nocturnal could be to avoid predators. This theory can explain the differences between megabats and microbats rather well. After flight was developed, several bats started to feed on fruits and nectar. These bats developed better sight and smell. Other bats kept feeding on insects and developed echolocation (Simmons et al., 2008; Speakman, 2001).

A different theory suggests the initial function of echolocation was communication between bats. Here, flight would have developed after echolocation. Primitive echolocation found in megabats such as clicking of the tongue supports this theory. Under this theory, megabats would have lost their echolocation abilities later (Speakman, 2001).

Lastly, there are theories that suggest bats relied on sounds to locate prey from the very beginning. These would be passive sounds from the prey itself. Hence their hearing improved over time and it was only a small step to start using the reflection of their own sounds to better locate preys. If an animal uses sounds made by a different animal, we call it passive mode hearing. If they use the reflection of their own sounds, we call it active

mode hearing. Megabats either split off before the transition from passive to active mode, or they lost their echolocation later (Speakman, 2001).

2.4.5 Classification of echolocation

Parameters of echolocation

A bat call has three important parameters: frequency, bandwidth and length. Each parameter can be varied resulting in different advantages and disadvantages (Denzinger and Schnitzler, 2013). These parameters are graphically shown in Figure 2.5




Figure 2.5 Different parameters of a bat sound.

As mentioned before, frequency ranges between 20 and 120 kHz with most calls being around 50 kHz. High frequencies create a high resolution of the environment. However, these frequencies also get absorbed better into the air, meaning the range becomes smaller (Lawrence and Simmons, 1982).

Bandwidth is the range of different frequencies present in a signal. A high bandwidth will make it easier to determine direction, size and distance of an object. However, it does not allow bats to use the Doppler effect or to tune their hearing to specific frequencies. As mentioned before, this is important in high-duty mode (Denzinger and Schnitzler, 2013).

Lastly, length (signal time) is important. A long signal will increase the signal to noise ratio which makes it easier to detect preys. However, echoes and calls can start to overlap, which is difficult to deal with (forward masking). Only long calls can detect the wings of insects (Denzinger and Schnitzler, 2013).

Guilds

Bats show some flexibility in adapting these parameters depending on the situation. Even so, most bats prefer sticking to a certain type of signal and have evolved to use this signal very effectively. Based upon this, they can be classified in different groups. One way to define these groups is in terms of guilds. A guild was defined by Root (1967) as a group of species that exploits the same class of environmental resources in a similar way. If animals

feed in a similar way, they belong to the same guild. We will discuss the five guilds from Denzinger and Schnitzler (2013), which are also used in the field guide of Dietz and Kiefer (2017). As an illustration, Figure 2.6 shows the signal of a representative of each guild.

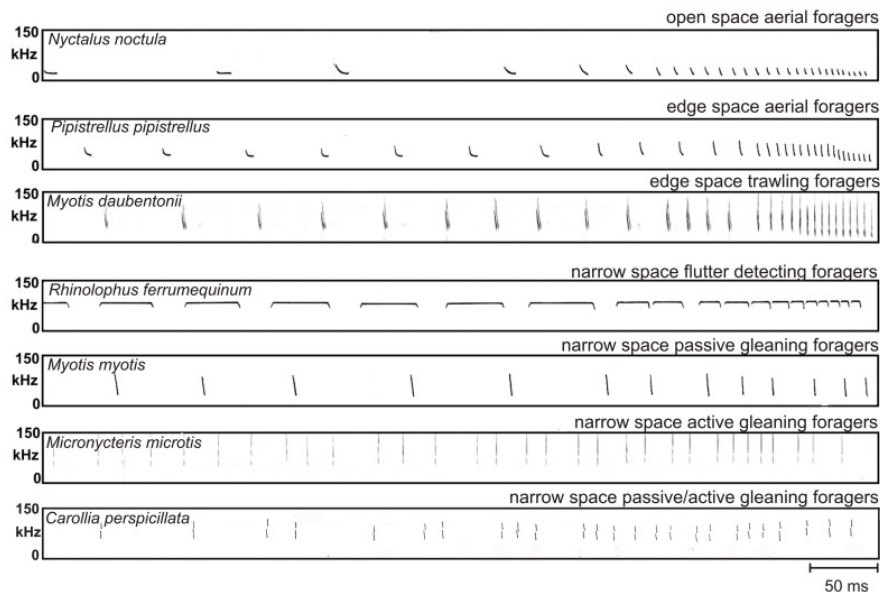


Figure 2.6 Signals of a representative of each guild. Source: Denzinger and Schnitzler (2013).

The first guild are the open-space aerial foragers. These bats hunt in open spaces and hence face a significant problem. Preys will be distributed over large areas which makes them difficult to locate. Hence, these bats will need long-range echolocation, which results in long calls at low frequencies (around 30 kHz) with a narrow bandwidth. Their morphology is adapted to this as well, with characteristic narrow wings to trade off extra speed for lower agility. Typical examples of these bats are *Nyctalus* and *Eptesicus* species.

The second guild are edge-space aerial foragers. These bats hunt in edge spaces, which are line-shaped structures such as hedges or the treeline of a forest. They face a challenge because there is a constant background of the edge structure (backwards masking). To counter this, these bats use mixed signals. They have a shallowly modulated narrow-band component which is either preceded or followed by a broadband, steeply downward frequency-modulated component. The first one is used for insect detection, the second one to fill in the surroundings (to avoid collisions). The first component (shallow one) is species specific and can be used for determination. The second component (steep one) is more general for this group. Typical members of this group are *Pipistrellus* species and *Barbastella* species.

Next are the edge-space trawling foragers. Trawling is a word for fishing with a net and this is reasonably accurate for this group. Trawling species will fish for signals by flying low above water surfaces while emitting calls. The smooth water will make the calls echo away from the bat except for the ones hitting the surface perpendicular. This is used to determine

the height above the water surface. The bat can easily detect objects on or above the water surface. Their signals are usually similar to the second guild because the shores work as an edge-space. A key difference is that trawlers tend to put a shallow modulated component in between two steeply modulated components, which makes it possible to distinguish between the two. However, when trawlers hunt above lakes, their signals are more similar to open-space aerial foragers. Examples of this guild are *Myotis* species such as *Myotis daubentonii* (Daubenton's bat).

The fourth guild are the narrow space flutter detecting foragers. Narrow space is sometimes changed to confined space instead. A confined space refers to an area which creates large amounts of echoes from all sides. The middle of a forest is a simple example, but urban areas can also serve as confined spaces. Backwards masking will be a significant problem here. To overcome this, these bats tune their calls to specifically detect the beating of insect wings. Insect wings are in constant motion which will create glints (rapidly changing echoes). These bats will have long calls at a constant high frequency. Morphologically they will have short, broad wings to increase their agility. An example of this guild are the *Rhinopholus* species.

Lastly, there is the guild of the narrow space passive gleaning foragers. Gleaning is a synonym for collecting. These bats will 'collect' insects from a surface. Rather uniquely, these bats do not use echolocation to locate their prey. Instead, they listen to the sounds of the insect itself (passive) along with sight and smell. As mentioned before, there is a theory that suggests that all bats were originally passive gleaners. Echolocation can be present for spatial orientation, however. In that case, their signals are similar to the first component of open-space aerial foragers. Examples of these bats are *Bechstein* species and *Plectocus* species. Bats that use a similar collecting technique with echolocation instead are called narrow space *active* gleaning foragers, but this is a very rare guild.

Finally, there is one more guild that is often ignored. Some frugivorous and nectarivorous bats use smell as their primary sense, but use echolocation to find the exact position of fruit or nectar once they are close. This guild is often called narrow space active/passive gleaning foragers, because they use both active methods (echolocation) and passive methods (smell). However, bats like this do not appear in moderate climates, so they will not be relevant to our research.

CHAPTER 3

CONSERVATION AND **MONITORING OF BATS**

In this chapter we will focus on the conservation of bats. First, we discuss the threats that bats face. This is followed by a section about possible conservation measures. Next, we delve into legislation to conserve bats, with a primary focus on the European Union and Belgium. Non-governmental organizations are also discussed briefly. Then, the current state of bats is discussed, starting worldwide and then zooming in on the European Union and Belgium. Lastly, possible monitoring methods are discussed.

3.1 Threats

Bats face many threats. Not all of them are unique to bats though. Habitat destruction for example is a threat to all animals, not just bats. In this section, both general threats and specific threats will be discussed. However, it is important to point out that even if a threat is general, it does not always hit species the same way. Bats in particular have certain properties that make them vulnerable. Most importantly, bats have a very low birth rate. Females usually only get one juvenile per year. Furthermore, it can take several years before a juvenile becomes sexually mature and starts to reproduce. This means that population growth is low, which makes recovery from events more difficult. Under normal circumstances, this low population growth is compensated by an exceptionally long lifespan. Wilkinson and South (2002) reports individuals of five bat species that lived for at least 30 years. This is about 3.5 times longer than the lifespan of a non-flying placental mammal of the same size. Additionally, bats have advanced systems to control time of birth. Some species store sperm during winter or freeze development of the embryo in the initial stages. Births happen in a narrow window around May or June in temperate climates. Normally this is an advantage, since food supply is optimal around this point. But if there is a disturbance during this narrow window, this can have disastrous consequences for a population (Dietz and Kiefer, 2017).

The first threat bats face are diseases and parasites. The most important example is the white-nose syndrome (WNS), a fungal growth. In the United States, this is monitored by the United States Geological Survey agency (USGS). They estimate that since 2008, millions of bats have died in the US and Canada from WNS. The disease was first seen in North America around 2006 and later on in Europe as well. However, in Europe there were no casualties. Current theories suggest the disease developed in Europe and hence European bats developed immunity against it. American bats did not have this evolutionary immunity, which caused mass casualties. Whether increased globalization caused WNS to spread from Europe to America is not clear (Fenton, 2012). Recent genetics research by Palmer et al. (2018) has found that the alternate excision repair pathway in the fungus is not functional. This means that the fungus can be killed easily by ultraviolet light or certain DNA alkylating agents such as methyl methanesulfonate. In the future, this could lead to a potential treatment of WNS.

Bats also contain a wide variety of other pathogens. Usually these do not pose a threat to the bat itself, but they can infect humans. Indirectly, this poses a threat to bats since they could be killed on purpose to prevent certain diseases from spreading. Viruses like *Ebola* have been reported in bats, as well as the *Corona* virus, which causes severe acute respiratory syndrome (SARS) in humans. The *Henipa* virus can be found as well. Lastly, there are also five variants of the *Lyssa* virus found in bats: EBLV 1, EBLV 2, WCBV, BBLV and LLEBV. *Lyssa* is known to cause rabies in humans, but so far, there are only three confirmed and two unconfirmed human deaths of bat rabies, all of which related to EBLV 1 and 2. Even so, it is recommended to wear gloves and use a vaccine when handling bats (Dietz and Kiefer, 2017).

Bats are also vulnerable to extreme temperatures. The most notable cases are the heat-waves in Queensland, Australia. In 2014, an estimated 45 500 flying foxes, roughly half of the population, died in a single day when temperatures reached up to 44.6 degrees Celsius (Welbergen et al., 2014). Another day in 2018 killed 23 000 individuals (Kim and Stephen, 2018).

Changes in natural predators are quite limited. Only owls, snakes, hawks, minks and racoons can catch a bat and eat it. However, house cats are a big threat to bats. Cats find a roost, catch bats as they emerge, and then play with their prey. They never eat a bat. The Bat Conservation Trust in the UK reports that 30% of rescued bats were attacked by house cats. Other research in Italy has found a very similar number (Ancillotto et al., 2013).

Apart from these somewhat natural causes, there are many direct anthropogenic threats. Roads and railroads are a common problem. Echolocation only works over short distances, so bats will never detect a fast moving train or car in time. In Greece, an entire population

was exterminated because a road crossed the route between their hunting grounds and roosting cave. Some species do not cross roads at all and fly around them or get stuck in one area. It is estimated that roads kill about the same number of bats as house cats, but exact numbers are difficult to estimate (Dietz and Kiefer, 2017).

Pesticides and chemicals can also pose a threat. Chlorinated components (HCHCs, DDTs, PCBs, CHLs,...) can accumulate through the food chain. Therefore, the highest levels are found in birds and bats (Senthilkumar et al., 2001). Pesticides can also have an indirect effect. If a pesticide is used to kill insects, the food supply of bats is reduced. Apart from pesticides, land transformation can also reduce insect populations (Matteson et al., 2013).

Changes in housing also affect bats. Many bats live in attics or other crevices in houses, but due to renovations, demolishing and better isolation, these habitats are destroyed. The most vulnerable of these habitats are winter resorts. These are places that are used by thousands of bats to hibernate. A famous example are the basalt mines of Mayen in western Germany. These old mines house between thirty and fifty thousand bats of sixteen different species according to the German Federal Agency for Nature Conservation (Bundesamt für Naturschutz). Around five million euro was spend on a conservation project. This project included extensive surveys of the populations and restoration of the mines. Entrances were barred with gates or sealed off and hazardous corridors were reinforced (Dietz and Kiefer, 2017).

Disturbance can be a problem as well. During the day bats go into a lethargic state and some bats hibernate in northern climates (others migrate). Both are vital to energy conservation. If bats are disturbed during lethargy or hibernation, it can have a significant effect on their health (Dietz and Kiefer, 2017). Disturbance can also come in the form of light. Apart from other effects, artificial light can delay the emergence from a roost, which gives bats less time to hunt (Stone et al., 2015).

Poaching can also pose a threat. In some South-American cultures, the blood of bats is believed to have healing properties. Lizarro et al. (2010) studied four major cities in Bolivia (Cochabamba, La Paz, Santa Cruz and Oruro) and found that over 3 000 bats were sold per month. But bats can also be hunted as a source of food. Goodman (2006) reports that in Madagascar, the *Hipposideros commersoni* is hunted during periods of food shortage.

Bats are also threatened by wind turbines. Turbines can cause barotrauma (lung rupture due to high pressure changes) or bats can hit the rotors. Exact statistics are limited, but wind turbines are estimated to be a bigger threat than cats and traffic (Baerwald et al., 2008; Dietz and Kiefer, 2017).

3.2 Conservation measures

There are several measures that can be taken to protect bats. A few of them will be discussed here.

Firstly, there is the conservation of known bat homes. If a cave is known to be a winter home for bats, cave gates can be placed to prevent people from entering and disturbing the bats. Bridges and houses can be checked for the presence of bats before being demolished or renovated. Of course, this is not always possible. Often demolishing is needed for safety reasons (Dietz and Kiefer, 2017).

If conservation is not possible, an artificial home can be constructed. These are called bat houses. They are commonly made from wood and provide space for roosting. The largest one was built by the University of Florida. They report a population between 450 000 and 500 000 bats (Florida Museum of Natural History, University of Florida, 2018). For the effectiveness of artificial bat houses we refer to Mering and Chambers (2014).

In Flanders, Natuurpunt and INBO measured the temperature and humidity of old ruins and bunkers around the city of Antwerp. Coupled with data about the preferences of different bat species and a heating model, management options can be evaluated. Usually management options are very simple measures such as the blocking of one or more openings (Vanheuverbeke et al., 2018).

3.3 Legislation

3.3.1 European Union

The main legislation in the European Union related to bats is the Habitats Directive (Council of European Union, 1992). This directive protects certain species of plants and animals. In annex IV, we find a list of protected species that mentions: microchiroptera: all species. The only other mammals with a similar high-level protection are the Cetacea (whales and dolphins). Other mammals are protected at the level of a family, genus or species.

Article 12 of this directive prohibits the following: deliberate capture or killing in the wild; deliberate disturbance (particularly during the period of breeding, rearing, hibernation and migration); deliberate destruction or taking of eggs from the wild; deterioration or destruction of breeding sites or resting places; keeping, transport and sale or exchange, and offering for sale or exchange, of specimens taken from the wild.

However, article 16 states that deviations from article 12 are allowed as long as there is no satisfactory alternative and the population can be maintained at a favorable conservation status in their natural range. The following reasons allow deviations: protecting wild fauna and flora and conserving natural habitats; prevent serious damage; for public health, public safety or public interest; research and educational purposes or re-introduction of species; limited taking in numbers specified by the competent national authorities (mainly related to monitoring). Article 16 was cited by the Commission during a dispute between the EU and the Church of England. The spokesperson for the Church claimed that church artifacts and wall paintings were being damaged by bats and the Habitat Directive did not allow them to do anything about this. In response to this, the Commission pointed out article 16 (Soady, 2013). However, the tension remained and several bills were introduced later to limit protection of bats in churches (Drake, 2015; Davies, 2016).

Lastly, article 17 states that all member states must report every six years on protected species. This report must include which conservation measures were taken and the effectiveness of them as well as the current status of the species. These results are published within two years in a report that is made publicly available. For the conservation status there are four classes: favorable, unfavorable-inadequate, unfavorable-bad and unknown. The trend is reported as: improving, stable, deteriorating or unknown. For meaningful comparison, the EU is divided into fourteen biogeographical regions. Every country needs to report on these regions separately. The regions are shown in Figure 3.1.

Since the adoption of article 17, there have been two reports. The first one published in 2007 for the period 2001-2006 and the second one published in 2014 for the period 2007-2012. The report for the period 2013-2018 is expected in 2020. The period 1992-2000 was used for the implementation of the directive and does not have any formal reporting.

3.3.2 Belgium

In Belgium, the protection of bats falls under environment and nature, which is a regional jurisdiction. This means that the three regions (Flanders, Brussels and Wallonia) work separately. Technically Belgium needs to report as a whole, but because Flanders and Wallonia are in different biogeographical regions, they can report separately. Flanders reports on the Atlantic region, along with Brussels. Wallonia reports on the continental region.

In Flanders, bat protection falls under the Department of Environment and Spatial Development. Within this department, the main relevant division is the Research Institute for Nature and Forest (INBO). They are responsible for the scientific research and monitoring of bat populations.

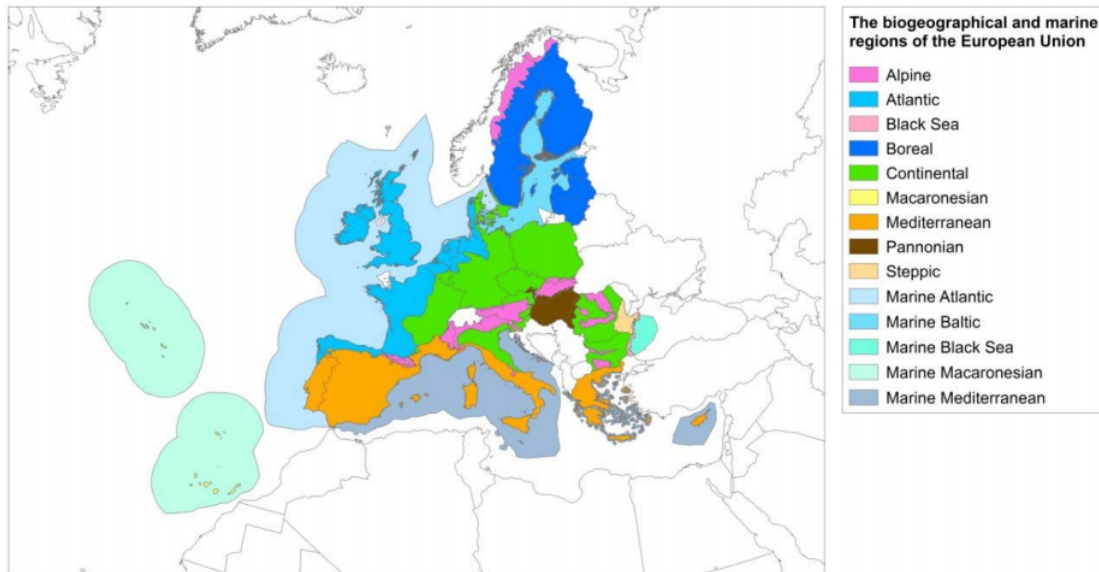


Figure 3.1 Biogeographical regions within the EU for the period 2007 – 2012. Note that Croatia is not yet included in the regions because they only joined the EU in 2013. The regions for 2013 – 2018 are given in the next report, which comes out in 2020. Source: Eionet, European Environment Agency (2013)

In Brussels, bats fall under Environment Brussels. Due to the small size of this region, actual monitoring of bats is very limited here.

In Wallonia, bats fall under the operational department for agriculture, natural resources and environment (D GARNE).

3.4 Non-governmental organizations

Various non-governmental groups also try to protect bats and raise awareness. There are international nonprofit organizations such as Bat Conservation International¹. This organization works on general conservation, education and research. Other organizations focus on more specific issues, such as Lubee Bat Conservancy² that works on the conservation of plants for fruit bats.

There are local organizations as well, such as the National Wildlife Federation (US)³, Florida Bat Conservancy⁴, BatLife Europe⁵ and Natuurpunt (in Flanders and Brussels)⁶. These organizations often work together, as they do on the international bat night held in the last weekend of August. During this weekend, presentations and bat walks are held in over thirty countries. In Flanders and Brussels this is organized by Natuurpunt.

¹<http://www.batcon.org/>

²<https://www.lubee.org/>

³<https://www.nwf.org/>

⁴<https://www.floridabats.org/>

⁵<https://www.batlife-europe.info/>

⁶<https://www.natuurpunt.be/>

Table 3.1 Overview of the current IUCN ratings of bats and birds (retrieved January 2019).

IUCN rating	Number of bat species	Percentage of bat species	Number of bird species	Percentage of bird species
Extinct (E)	5	4.02	156	1.40
Extinct In The Wild (EW)	0	0	5	0.04
Critically Endangered (CR)	23	1.85	224	2.01
Endangered (EN)	56	4.50	469	4.22
Vulnerable (VU)	106	8.52	799	7.18
Near Threatened (NT)	82	6.59	1012	9.10
Least Concern (LC)	745	59.89	8405	75.54
Data Deficient (DD)	227	18.25	56	0.5
Total	1244	100	11126	100

Apart from these organizations, there are also organizations that do not have the main goal of protecting bats or nature, but still do a lot of conservation work or research. Universities are an example of this. Universities mainly do research, but actual conservation is possible as well. The University of Florida and their artificial bat house was already mentioned.

Lastly, individuals can take action as well. Many organizations have webpages about first aid for wounded bats and the construction of bat houses.

3.5 Current status

3.5.1 Worldwide

The International Union for Conservation of Nature (IUCN) monitors the state of animals and plants. This is published in the Red List. There are nine categories: Not Evaluated (NE), Data Deficient (DD), Least Concern (LC), Near Threatened (NT), Vulnerable (VU), Endangered (EN), Critically Endangered (CE), Extinct in the Wild (EW) and Extinct (E). For the last seven categories, a list of criteria is evaluated. If the species does not qualify for these criteria, it is labeled LC. If it does not qualify, but is likely to do so in the future, it is labeled NT. If it qualifies, it is labeled VU, EN or CE depending on how high the risk is. Animals in these three categories are considered to be threatened with extinction. The last two categories are reserved for animals that are extinct in the wild already (EW) and animals that are completely extinct (E).

The current status of bat species according to the IUCN is shown in Table 3.1. As a comparison, birds are also shown. We can see that roughly 15% of bat species are at risk for extinction (CE, EN and VU). Birds show a very similar number. In terms of already extinct species, bats are clearly higher with around 4% compared to only 1.4% for birds. However, the biggest difference is in the Data Deficiency category. Relatively speaking, for every bird species without sufficient data, there are 36 bat species in the same situation. This clearly illustrates that a large challenge lies in the monitoring of bats. There are more bat species without sufficient data than there are actual threatened bat species.

3.5.2 European Union

As mentioned before (Section 3.3.1), article 17 of the habitat directive requires member states to report to the EU every six years. The specific entity they report to is the European Environment Agency (EEA). This agency then has two years to compile the information. Detailed information is published on Eionet (Eionet, European Environment Agency, 2014). Both assessments for habitats and species are published. We will only focus on the species assessments.

A species assessment needs to be done per member state for all species that are present in a specific biogeographical region. This means that if a country has two biogeographical regions and a certain species is present in both, two separate assessments must be made. A full species assessment contains three elements: range, population and suitable habitat. The population part contains an estimate of the number of individuals present in the area as well as the trend. A trend can be stable, increasing, decreasing or unknown. There is also a favorable reference value, which is an estimate of a 'good' number of individuals to be present in this specific area. Based on the population compared to the reference value and the current trend, the status can be favorable ('green'), unfavorable-inadequate ('amber'), unfavorable-bad ('red') or unknown ('gray'). Range and suitable habitat follow a similar procedure to get to a status. These three statuses are then combined into a single overall assessment.

For every member state, there is a national summary which publishes several graphs based on the overall assessments. In general, this makes things easier to analyze. However, in this summary, bats are grouped together with other mammals, which makes it difficult to draw a conclusion for bats alone.

3.5.3 Belgium

The status of bats in Belgium is shown in Figure 3.2. This figure is created based on the data on overall assessments for Belgium in the periods 2001-2006 and 2007-2012 in the regions continental (CON) and Atlantic (ATL). As is clear from the figure, there are no changes between the two periods. Furthermore, the status is significantly better in the Atlantic region. It is also notable that a significant number of species have an unknown status. Clearly, the monitoring programs have room for improvement.

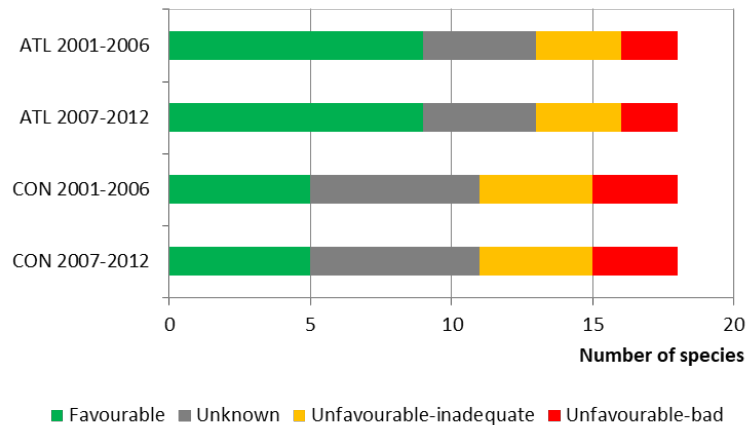


Figure 3.2 Overall assessments on Bats in Belgium. ATL: Atlantic region (Flanders and Brussels), CON: Continental region (Wallonia). Source: Eionet, European Environment Agency (2014).

3.6 Monitoring

A big problem in bat conservation is the monitoring of bat populations. It is difficult to determine whether measures have any effect unless the populations are carefully monitored and mapped so that changes can be detected. This is rather challenging since, as discussed, bats are nocturnal flying animals, which makes them difficult to track. Nevertheless, several methods have been developed.

3.6.1 Non-audio methods

We will briefly discuss a few common non-audio monitoring techniques: marking, transponders, telemetry, light traps and winter counting.

Marking and transponders are two methods that are related to one another. The bat is caught with a net and then a device is attached to it. In the case of marking this is a simple iron ring with a serial number on it. Transponders use a microchip that can be read by a scanner. Transponders have the advantage that the animal only needs to be caught once. By attaching a detector to a known roosting area, incoming bats can be scanned automatically. These two methods are commonly used to map migration patterns (Dietz and Kiefer, 2017).

Telemetry is related to marking as well. However, instead of a microchip or ring, a transmitter is used. With a receiver, the location of the bat can be determined. Transmitters are attached with a collar with a degradable link or with medical glue. In the case of glue, they stay on for about nine days (O'Mara et al., 2014). An important limitation here is that the weight of the transmitter cannot be too high. In theory, the guideline is less than 5% of the

total bodyweight. However, research by O'Mara et al. (2014) reports that 54% of the studies do not report individual body mass or forearm length (which is a proxy for body mass). Of the studies that did, 51% followed the guideline, while 47% exceeded the guideline but stayed below 10%.

Light traps are another method. This technique uses two light curtains of infrared rays and detectors placed at the entrance of a roosting area. When a bat flies through this opening, the signal is interrupted, which allows us to detect a passing bat. The order of the two signals determines whether the bat comes in or goes out. This can be expanded with a camera trap that snaps a picture of the bat for determination (Dietz and Kiefer, 2017).

Lastly, there are counting techniques. Mainly winter counting is important here. Because most bats hibernate, it is possible to simply go inside a winter home and count the hibernating bats. Determination can be difficult because most features are hidden. However, bat species are known to have a different preference in terms of temperature (Webb et al., 1996). This can sometimes help with determination. The bats emerging from a roost can also be counted. The scale of these studies can be quite large. The National Bat Monitoring Programme of Great Britain took place between 1997 and 2012 and had 3 500 volunteers counting individuals of ten bat species. With this data, the evolution of the bat populations could be followed (Barlow et al., 2015). In Wales, winter counting was used to estimate the size and evolution of the lesser horseshoe bat population (Warren and Witter, 2002).

3.6.2 Audio methods

Apart from non-audio methods, there are also audio methods. Audio methods use recordings of echolocation calls and try to identify species based on this audio. Different species will have different calls. These techniques are part of the field of ecoacoustics that is discussed in Chapter 4.

CHAPTER 4

ECOACOUSTICS

The most common methods to monitor bats are audio methods. Audio methods are part of the field of ecoacoustics. In this chapter, a brief history of ecoacoustics will be given along with the most important applications. Then, two key hypotheses within ecoacoustics will be explained. Lastly, we will focus on bats. First, the hardware and software of bat recorders will be discussed, followed by a discussion of the technical problems and challenges related to the recording of bat sounds.

4.1 What is ecoacoustics?

Ecoacoustics is a recent field of research that uses environmental sounds for ecological purposes. Common synonyms are soundscape studies and acoustic ecology. Soundscape refers to the World Soundscape Project, which is usually considered the first research team that studied ecoacoustics. Their first study was done in Vancouver in 1973 (Schafer et al., 1973). This study mainly focused on the sources of urban noise and the reaction of humans to this noise. Later on, ecoacoustics was expanded to other applications. We will focus on ecoacoustics within the context of ecological monitoring.

4.2 Acoustic niche hypothesis and acoustic adaptation hypothesis

There are two important hypotheses in ecoacoustics: the acoustic niche hypothesis (ANH) and the acoustic adaptation hypothesis (AAH).

ANH was first formulated by Krause (1993). He observed during a monitoring program that 'The bird, mammal and frog vocalizations we recorded all seemed to fit neatly into their respective niches'. In other words, there is a partitioning of the acoustic space. Different species will make different sounds. This sounds trivial, but it is an important underlying assumption. If different animals made the same sounds, it would not be possible to tell them apart using those sounds.

AAH says that the habitat itself will shape the sounds created by animals to maximize propagation. Sound will propagate differently in a wide open field compared to a forest. Therefore, animals who live in a forest will develop sounds that work best in a forest. It also means that the same species in a different habitat can make a different sound. For bats this will be very important to keep in mind. This idea was put forward in multiple studies. A review of these studies can be found in Ey and Fischer (2009). For every study, it is important to think about these two hypotheses. If two sounds are different, this could be due to the species (ANH) or the environment (AAH).

4.3 Applications of ecoacoustics

Ecoacoustics is commonly used in monitoring problems that do not provide easy access for other methods. Underwater animals for example are difficult to monitor due to their habitat. Širović and Hildebrand (2011) and McCauley et al. (2018) studied whales using underwater acoustic sensors. Other applications are in fields where sound is very strongly linked to species. Bird studies are a classic example of this. Farina et al. (2011) suggested a methodology for long-term bird monitoring using ecoacoustics.

But applications can go far beyond the study of individual animals. Research by Krause and Farina (2016) found a relationship between the sounds of animals and climate change. Small changes to fundamental aspects of the environment such as temperature, humidity, pH or vegetation cover can result in different sounds. Through time series analysis, the impact of climate change can be estimated.

Ecoacoustics can be used in various disciplines that are closely related. Some studies focus more on biodiversity, while others might focus more on habitat or conservation. For an overview of the research in these various fields, we refer to Sueur and Farina (2015).

Ecoacoustics is used widely in bat studies for the reasons mentioned above. A lot of monitoring methods, such as camera traps, do not work well on nocturnal flying animals. Furthermore, the sound a bat makes shows a strong link with the species. However, care must be taken. Bat calls are sometimes compared with bird calls, but there is a very important difference. Birds communicate their identity to the receiver and therefore, it is evolutionary advantageous for their calls to be species specific. Bats on the other hand use their calls to interact with their environment, so there is no pressure for different bats in similar environments to have different calls (Barclay, 1999).

4.4 Bat detectors

A bat detector is a device that uses the echolocating calls to detect and/or classify bats. Simple bat detectors convert ultrasound to audible frequencies. As mentioned before, bats emit calls roughly in the range 20-120 kHz. It is possible that the upper limit is actually higher, however, bat detectors usually do not measure above 120 kHz. Furthermore, high frequencies are absorbed more quickly into the air and are hence more difficult to detect. When the frequency is changed or modulated people can learn to recognize specific sounds and associate them with certain bat species. Simple bat detectors use this manual method. More advanced bat detectors record the calls and store them so they can be analyzed with a computer later on (Dietz and Kiefer, 2017). Figure 4.1 shows a simple handheld bat detector. Figure 4.2 shows a more advanced passive detector that can be left out in the field. Table 4.1 shows a summary of the main types of bat detectors with their advantages and disadvantages.



Figure 4.1 Example of a heterodyne bat detector. (Magenta bat 4 bat detector, Magenta Electronics Ltd.)



Figure 4.2 Example of a full spectrum passive bat detector. (Anabat Swift bat detector, Titley Scientific.)

4.4.1 Heterodyne detectors

The simplest type of bat detector is a heterodyne. ‘Hetero’ means different and ‘dyne’ means power. Different refers to the two different frequencies used. First, there is the external frequency from the bat itself. Let us say this is a call of about 50 kHz. Secondly, there is an internal frequency that is set by the user. Let us assume we put this at 45 kHz. The bat detector will now play a constant tone at 45 kHz. This creates interference. There is constructive interference, which will add up the frequencies to 95 kHz. Then there is also destructive interference, which subtracts the frequencies. This gives us 5 kHz, a frequency that is audible by humans. Interference is illustrated in Figure 4.3. Whenever a bat flies by, clicking noises can be heard from the heterodyne (Cutnell and Johnson, 1997; Boonman et al., 2000). When this device is used, the standard method is to set it to 45 kHz and then manually adjust the dial until the sound is most audible. Actual classification is rather difficult as only the average frequency is available. Because of this, heterodynes are usually used in the field or for bat walks. However, classification is possible. In a study in Cuba a simple heterodyne was enough to classify five bat species (Moreno et al., 2016).




Figure 4.3 Illustration of constructive and destructive interference. This principle is used in a heterodyne. Source: HyperPhysics (R. Nave)

Heterodynes can be expanded in several ways. A common expansion is a scanning circuit. This adds a microprocessor that automatically adjusts the internal frequency in fixed steps. When a bat call is detected, it stops. This avoids constant manual adjustments. It is also possible to use a comb generator. Comb generators create harmonic inputs, which results in multiple internal frequencies being present in the same signal, effectively tuning it to a much wider range of frequencies (Dietz and Kiefer, 2017).

4.4.2 Frequency division detectors

Frequency division (FD) bat detectors work differently compared to a heterodyne, but have the same goal. Just like a heterodyne, they will try to generate an audible signal. An FD device does not transform the real signal, but generates a new signal based on it. To that end, it counts the number of pulses within the original signal. Then, it will send out a pulse for each n pulses it detects, where n is usually 10. If a bat call has a constant frequency of 50 kHz or 50 000 pulses per second, then for every ten pulses the detector counts, it sends out one pulse, resulting in a new signal of 5 kHz. This signal is audible. Just like heterodynes, FDs are mostly used in the field (Boonman et al., 2000).

4.4.3 Digital detectors

There are also digital detectors. These are different from the previous type, because they actually record the sounds. A digital detector will always have a microphone combined with a converter. The sound is recorded with the microphone and transformed into an analog signal, which is then converted further into a digital signal (Boonman et al., 2000).

The microphone is not a normal microphone as might be found on a smartphone. This is because the sampling rate will not be high enough. All sounds are waves. The only thing a microphone does is measure the height of this wave at a number of points in time. The number of measurements per second is called the sampling rate. The higher the frequency of the waves, the higher the sampling rate needs to be to accurately reconstruct them. The Nyquist-Shannon sampling theorem states that the minimum sampling rate is twice the frequency of the sound that is recorded (Shannon, 1949). A 50 kHz signal requires at least 100 kHz sampling rate. Normal microphones do not need to go much above 40 kHz, because audible sounds will not go above 20 kHz. But bat calls can go up to 120 kHz, which means a special microphone is necessary that can measure six times faster than a normal microphone. In practice, the practical limit of a microphone can be lower than the theoretical limit based on the Nyquist-Shannon theorem. Most professional microphones will mention their practical limit in the technical specifications. Usually this is about 120 kHz (Boonman et al., 2000).

Once the sound is converted into a digital format, there are two main options: playback and storage. Based upon this, there are four subtypes: time expansion, full spectrum, zero-crossing and time domain signal coding.

Time expansion detectors

Time expansion (TE) detectors use playback. These are digital detectors with a very low memory that only serves as working memory. The signal is stored temporarily and then played back at a slower speed. Usually a factor ten is used (Boonman et al., 2000).

This may sound like a FD detector, but there are two important differences. First, no new signal is generated, the original signal is simply played at a slower speed. Secondly, TE detectors have a delay, because they record the full call first and then play it back. FD detectors work in real time (Boonman et al., 2000).

Full spectrum detectors

It is possible to simply store the full digitally converted signal. These devices are called full spectrum bat detectors. Active full spectrum bat detectors are handheld devices that display a spectrogram. Passive devices have higher storage and are left in the field. A big problem here is that due to the high frequency sampling rates, the files become quite big. A single night can easily lead to several gigabytes of data (Boonman et al., 2000).

Zero-crossing analyzer

The zero-crossing analyzer (ZCA) is designed to combat the memory problem. Instead of storing the whole spectrum, a ZCA only stores the zero crossings. Zero crossings are the points where the spectrum intersects the x-axis. In other words, we only store information about the frequency, not the amplitude. While this saves memory space and reduces power consumption, it is not ideal. Amplitude can also contain information about the bat species. Furthermore, if two different sounds are present at the same time, the ZCA method does not work properly. Either it will only store one of them, or it will mix both sounds together, which makes identification difficult (Boonman et al., 2000).

Time domain signal coding

The time domain signal coding (TDSC) detector is an expansion of the zero-crossing analyzer. Instead of saving only the zero crossings, it also saves extra information about the signal between crossings. This information is extracted and mapped to a specific code. For example: the number of local minima and maxima between two crossings and the time between the crossings can be extracted. A simple table then transforms these parameters into a single code. If there are three local minima and the duration is between five and seven ms, this could correspond to code number seven. With this system, every set

Type	Advantages	Disadvantages	Price range (euro)
Heterodyne	-Cheap -Real-time	-Manual adjustment -Narrow bandwidth -No harmonics	30 – 4000
Frequency division	-Cheap -Real-time -High bandwidth	-Limited accuracy -No harmonics	80 – 4000
Time expansion	-Almost real-time -More information	-No storage -Expensive -No harmonics	1000 – 4000
Full spectrum	-Highest information -Harmonics detectable -Lowest interference	-High storage space -Power consumption -Expensive -Analysis time	1000 – 4000
Zero crossing	-Accurate	-High interference -Analysis time -No harmonics	500 – 4000

Table 4.1 Summary of bat detectors. Source: Brigham et al. (2004) and Boonman et al. (2000). Price ranges derived from the nhbs wildlife store (retrieved: 2019-03-05). Upper price range includes devices that have multiple detection methods.

of parameters is mapped to one specific code. This will transform the signal into a list of numbers, which can be analyzed (Chesmore, 2001).

Combinatorial devices

Many modern bat detectors can function as multiple types. For many of these detectors, the same hardware or very similar hardware is required. Full spectrum detectors for example commonly have software inside to convert the spectra to zero crossings. Some full spectrum detectors also have an audio jack which allows them to function as time expansion devices. It is also common to combine an FD detector with a heterodyne. But there are other expansions possible as well. Some devices have a GPS function to save the location of the recorded calls. In passive detectors, it is possible that an expansion is installed to let it operate remotely using SMS commands.

4.5 Commercial software for classification of bats

Digital signals can be analyzed using specific software. Several packages are available. As a general example, SonoBat¹ will be discussed to highlight the disadvantages of the available software.

Firstly, there is a localization problem. Different bats live in different areas, so a high performance will require the software to be trained for the specific region one is interested

¹<https://sonobat.com/>, accessed: 2019-03-03

in. SonoBat has fifteen versions. Thirteen versions are for regions in the US, one for Alaska and Canada and another one for the rest of the world. Furthermore, the international version does not actually classify bats, it only extracts certain parameters from the data that can be used to build a classifier. Only the US versions actually classify species. This makes the software almost unusable in most of the world.

Secondly, SonoBat is commercial software. The regular software cost \$1526 for the US versions and \$680 for the international version (January 2019). This is a significant hurdle to use the software. It also makes the exchange of results rather difficult. If a third party wants to repeat an analysis of audio data, they need to buy a license as well. Like most software packages, SonoBat uses an internal reference library and custom feature extraction which are not available to third parties due to copyright.

Thirdly, there is customization. Since the source code of SonoBat is not available to the user, there is no room for any kind of improvement. This is a severe limitation, because specific recordings are likely to have specific problems. For example, there might be a known source of ultrasound present near the recorder. When software is adaptable, a simple piece of code can be written to filter this out. If the source cannot be adapted, the code is likely to make the same mistakes over and over again.

4.6 Technical problems with bat detectors

Various technical problems can arise when a bat detector is used. Three main types are defined: problems inherent to recording sound, variation between different sensors and failure of the sensor called anomalies.

4.6.1 Recording

Regardless of the sensor, bat calls are almost never recorded perfectly. When the distance to the bat increases, the higher frequencies are absorbed first. As such, a part of the spectrum will sometimes be missing, making determination very difficult. Spectra are also not cleanly separated. Often, multiple bat calls are recorded at the same time and their spectra overlap significantly. Frequencies also change due to the Doppler effect. Lastly, a call is often recorded multiple times. First, the call itself, but then shortly after also the reflection from the ground and the echoes of the call. These echoes can overlap with the call itself making it difficult to extract a full call (Adams et al., 2012). Figure 4.4 shows an example of an overlap. The second pulse shows a lower intensity, which means this is most likely an echo caused by sound reflecting off the ground. But it could also be a second bat that is farther away. Figure 4.5 shows two examples of multiple calls that are too close

together to separate. There is no intensity difference here, so these are most likely two bats passing by at the same time.




Figure 4.4 Calls from a *Pipistrellus pipistrellus*. The left figure shows a normal call, the right figure shows an overlapping call. Calls are from domain Huizingen. Data provided by INBO.




Figure 4.5 Calls from a *Miniopterus schreibersii* (left) and a *Pipistrellus nathusii* (right). Data from Barataud (Barataud, 2015).

4.6.2 Variation in detectors

Different detectors might record the same call differently. The general microphone sensitivity could be different, but there could also be an influence of direction and/or frequency on the sensitivity of the microphone. The algorithm to start recording could be different as well. Many bat detectors do not record continually, but only start when they detect ultrasound. This detection algorithm can be different for different instruments (Adams et al., 2012).

4.6.3 Anomalies

Lastly, there are anomalies or sensor failures. In Dauwe et al. (2014) city noise is being analyzed. The authors define three types of anomalies related to sensor failure. Firstly, abrupt failure. This is a permanent deviating behavior of the sensor. Usual causes are broken microphones or water that makes contact with electronics. Secondly, incipient faults. These are small continuous faults. Sensor drift is the most common form. It is for example possible that a sensor will consistently overestimate the frequency of all sounds. To prevent this, instruments need to be calibrated periodically. Lastly, there are intermittent faults. Intermittent faults are temporary wrong measurements. Common causes of this are extremes in temperature, precipitation or other weather conditions. In general, the operator will check for anomalies and remove them before the data is analyzed.

4.7 Challenges with bat sounds

Sometimes, problems can arise due to the nature of bat calls and not because of the detector. In general, there are three main challenges with bat sounds: two different species can have similar calls (convergence), two individuals of the same species can have different calls (variation) and a different source of ultrasound could look like a bat call (interference).

4.7.1 Convergence

It is possible that two different bat species have very similar calls. This is not so common in birds, because bird calls play a role in communication and breeding, so it would be a disadvantage if two birds had the same call. But bats use echolocation for orientation and hunting, so it is entirely possible for two different species to have a similar call. In fact, it is even quite likely to happen. As mentioned before, bats have to take into account the range of their own hearing, the Doppler effect and the absorption of frequencies in air. As such, it is likely that two different bats in a similar environment would develop similar calls, because those calls are the most efficient in that specific environment. This can happen through convergent evolution, where species independently find a similar solution to the same problem, or because multiple species retain the calls of a common ancestor. This makes it challenging to draw general conclusions about bat sounds and species without considering local variation (Barclay, 1999).

4.7.2 Variation in calls

As mentioned before in the AAH, animals evolve different sounds for different environments. Bats take this several steps further. They do not have a single call, but rather a range of calls they use for different environments. In open land, bats can afford to use long pulses because they take a long time to return. In denser environments such as forests, bats tend to shorten their calls because the echoes return faster. The function also changes the pulses. Bats have different modes such as low-duty cycle and high-duty depending on whether they are mapping their environment or trying to catch an insect. When they are flying towards an insect, their calls can also change depending on the distance to their prey. Bats also have social calls to communicate with each other which tend to have a much lower frequency. Between bats of the same species there can also be differences based on properties such as age, sex and reproductive state. As such, it is rarely possible to tie one call to one specific species (Dietz and Kiefer, 2017). Figure 4.6 shows three calls, all from the same location and the same species (*Pipistrellus pipistrellus*). It is clear that even for bats of the same species in the same environment, there is some variation present. In

different environments, this difference is even larger. As an illustration, Figure 4.7 shows two calls from a *Myotis daubentonii*: one in an open region along a river (left), another one in a mountain range (right). The calls become much longer in the mountain range, most likely to increase the range and fill in more information about the environment.




Figure 4.6 Three calls from a *Pipistrellus pipistrellus*. Individual pulses were extracted from spectrograms. These calls are from the same day in the same location (Huizingen), but still show some variation. Data provided by INBO.




Figure 4.7 Calls from a *Myotis daubentonii*. The left figure shows this bat in Haute-Vienne, a department in the center of France along the Vienne river. The right figure shows the same bat species in the Alpes-de-Haute-Provence department. This department is located in the southeast of France and is part of the Alpes mountain range. Data from Barataud (Barataud, 2015).

4.7.3 Interference

Bat detectors record ultrasound, but bats are not the only source of ultrasounds. Insects, traffic, rain, pressing the button of the detector, clothes and many more things can be sources of ultrasound (Adams et al., 2012). Sometimes, interference can be very subtle. In research by Surlykke and Kalko (2008), researchers found that bats flying above low water recorded their peak frequency up to six decibels higher due to the reflection on the water surface. However, the main source of interference are jamming signals produced by tiger moths and hawkmoths (Corcoran and Moss, 2017). For an extensive discussion of tiger moths jamming bat echolocation, we refer to Corcoran et al. (2009).

CHAPTER 5

MACHINE LEARNING FOR DATA

ANALYSIS, CLASSIFICATION AND

CLUSTERING

In this chapter, we will discuss computational tools to automatically analyze, cluster and classify data. We will first discuss what machine learning is and what the different types are. Then, we will discuss unsupervised learning in detail. Lastly, we will talk about data transformation techniques and unbalanced datasets.

5.1 What is machine learning?

Machine learning refers to the scientific field that uses computers to perform specific tasks without explicitly programming them to do so (Chollet, 2017). The computer does not follow a specific set of instructions, but rather learns the instructions from the data itself. This is very different from classical explicit programming. A graphical representation of this new paradigm is given in Figure 5.1. In machine learning, after a specific model is selected, the model is trained using a training dataset. This dataset contains examples relevant to the problem at hand. Based on these examples, the model will search for rules that solve this problem (Chollet, 2017).

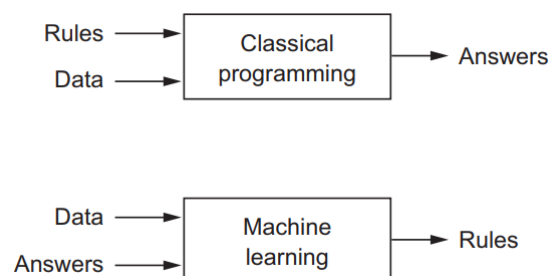


Figure 5.1 Illustration of the paradigm in machine learning versus regular programming (Chollet, 2017).

5.2 Types of machine learning

Based on the type of the training dataset, there are five types of machine learning: supervised, unsupervised, self-supervised, semi-supervised and reinforcement learning. Only supervised, unsupervised and semi-supervised will be discussed, as these are the most relevant to this work.

5.2.1 Supervised learning

Supervised methods are the most common by far. Here, the training dataset consists of labeled data, data where every sample has been given one or more labels. The goal of the model is to predict the output or labels solely based on the input, usually called the features. These features can be defined by the user or extracted by the model itself. Because a labeled dataset is already available, there is a way to get the labels without actually using the model. The true value of supervised learning lies not in solving new problems but solving problems much faster. It is possible to find the output manually, but the process is slow or expensive. Facial recognition is an example of supervised learning. Humans can recognize faces on pictures easily, but a computer can do it much faster (Chollet, 2017). One form of supervised learning is distance metric learning, which will be discussed in Section 5.4

5.2.2 Unsupervised learning

In unsupervised learning, the training data has no labels. In this case, supervised methods cannot be used, since performance cannot be measured and optimized. But not everything is lost in this case. We would still expect similar labels to show similar features, even if we do not have those labels available. For example, consider the problem to classify bat calls. We know that these signals belong to several species of bats. However, we do not know which call belongs to which bat. However, if we can define a metric to compute the similarity between two bat calls, we would expect that two calls of the same bat species would have high similarity and two calls of different bat species would have low similarity. Based on this, it is possible to try to separate the data into different groups (clusters). This technique is called clustering. Most unsupervised learning methods are not used for prediction, but rather for exploration, visualization, compression, denoising or understanding correlations between variables (Chollet, 2017).

5.2.3 Semi-supervised learning

Semi-supervised learning is a special class of machine learning that uses both unlabeled and labeled data. This is common in situations where the cost to label data dramatically outweighs the cost to collect data. Typical examples are data that require experts for labeling or data that requires physical experiments to label. Audio calls from bats can sometimes fall into this category since they generally need to be labeled by experts. However, this is not always the case. If an audio call is labeled at the level of a full call, that might not be useful if the model works with individual pulses. Semi-supervised learning is based on an important assumption: the continuity assumption. The continuity assumption states that points that are closer together are more likely to share the same label. Therefore if one datapoint has a known label, it is likely datapoints close to it share this label (Chapelle et al., 2006).

5.3 Clustering methods

Because the data is only partly labeled, we will mainly focus on unsupervised learning methods. A few of the most common clustering methods will be discussed. First a cluster needs to be formally defined. The general definition of a cluster is a group of datapoints that are more similar to each other than they are to datapoints in a different cluster (Estivill-Castro, 2002). Similarity can be defined in many different ways. Based upon how this is done, there are different clustering methods. Two important ones will be discussed: *K*-means clustering and self-organizing maps.

5.3.1 *K*-means clustering

K-means clustering is one of the most basic algorithms in clustering. It makes use of centroids. The centroid of a group of datapoints is defined as the center of mass of a set of feature vectors. This is the average value in each dimension. For example: the centroid of the points A (1, 3), B (1, 2), C (2, 2) is $(4/3, 7/3)$.

The concept of centroids can be used to define similarity in a single metric. This metric is the within cluster sum of squares (WCSS). Every cluster has a centroid. For every feature, the squared difference with the corresponding centroid is calculated and all of these numbers are added up. This is done for every datapoint and all these numbers are summed up. This gives the following equation:

$$WCSS = \sum_{i=1}^k \sum_{x \in S_i} \|x - \mu_i\|^2. \quad (5.1)$$

In this equation, different clusters are denoted as S_i and μ_i is the centroid of cluster i . x refers to a datapoint and there are k datapoints in total. Using this technique, the total variance within all the clusters combined can be described with a single number. For the same datapoints and number of clusters, a lower value indicates a better clustering, since the points are closer to their corresponding centroid. This metric is then minimized by changing which datapoints (x) belong to which cluster (S_i). This method is called the K -means clustering, sometimes referred to as Lloyd's algorithm (Lloyd, 1982; MacQueen, 1967).

In practice, this method works iteratively. An initial set of centroids is defined first. Then, every datapoint is assigned to the nearest centroid. Based on these new clusters, the centroids are calculated again. In general, this algorithm keeps going until the centroids stop moving around. This does not guarantee that the global optimum is reached (Hartigan and Wong, 1979). Depending on the initial conditions, it is possible to get stuck in a local optimum. As such, it is common to run the algorithm multiple times with different initial conditions.

5.3.2 Self-organizing maps

Self-organizing maps (SOM), also called Kohonen maps or self-organizing feature maps (SOFM) are a form of unsupervised learning. SOMs consist of neurons that each have a weight vector. These weight vectors have a dimension that is equal to the number of features in the dataset. SOMs use competitive learning, meaning each neuron will compete for the right to react to a specific subset of inputs. Maps are updated until the datapoints are optimally connected to the neurons (Kohonen, 1982).

Let us clarify this with an example. Consider four datapoints with two features that are measured: A (1, 2), B (2, 1), C (3, 4) and D (4, 3). We will try to classify these points using two neurons, each with their own weight vectors: W_1 and W_2 . Initially, the weight vectors are set randomly. Let us say the initial settings are: W_1 (0, 0) and W_2 (5, 5). This is illustrated in Figure 5.2a. Note that a weight vector always has the same number of dimensions as the input data and can be plotted like a hypothetical datapoint. Now, an iterative process starts. First, a random datapoint is selected, say datapoint C. The Euclidean distance from this point to each neuron is calculated. In this case, this would be 13 for W_1 and 3 for W_2 . The neuron with the smallest distance is called the best matching unit (BMU) or winning neuron. Competitive learning means there can only be




Figure 5.2 Illustration of a self-organizing map. Red dots are datapoints, blue triangles are weight vectors. The weight vectors move towards the datapoints. In the end setting, points A and B can be represented by weight vector W1 and points C and D can be represented by weight vector W2. Two clusters have formed.

one winning neuron. This winning neuron will now update its weight vector towards the datapoint according to the following equation:

$$W_u(s+1) = W_u(s) + \alpha(s) \cdot (D(t) - W_u(s)). \quad (5.2)$$

In this equation, s is the index of the timestep, which goes up with every iteration, $W_u(s)$ is the weight vector of the best matching unit at timestep s . $D(t)$ are the features of the randomly selected datapoint t . α is a learning coefficient that always lies between zero and one. When it is zero, weight vectors will never move around and nothing is learned. If the learning rate is one, the updated weight vector coincides with the selected datapoint. This would cause the neurons to jump around constantly between datapoints and the map would never converge. To find a balance between these two extremes, the learning rate is decreased monotonically after each iteration. We call this shrinkage. Shrinkage causes the map to converge as the algorithm runs. The easiest way to implement this is to pick an initial learning rate and a decay coefficient. If the decay coefficient is 0.9, then our learning rate at step $s+1$ is 0.9 times the learning rate at step s . More complicated functions are possible as well (Kohonen, 1982).

If this is applied, it is clear what will happen. If A or B is selected, the BMU will be W1. If C or D is selected, the BMU is W2. This will cause W1 to move towards A and B and W2 to move towards C and D (Figure 5.2b). Because all points have the same chance of being selected, the neurons will move to the exact midpoint between their two datapoints. The weight vectors will move less in each step because the learning rate decreases. In the final configuration, the average distance of every datapoint to the nearest neuron is as small as possible. In this case, it is 0.5. There are other configurations that have the same average distance in this case, but not a single configuration has a lower average distance. This

means that this is effectively a clustering method. Points that are close together, will have the same BMU, while points that are farther apart will have a different BMU. This can be expanded even further by connecting the neurons. This updated equation becomes:

$$W_v(s+1) = W_u(s) + \theta(u, v, s) \cdot \alpha(s) \cdot (D(t) - W_v(s)). \quad (5.3)$$

There are two differences here. Firstly, we consider the neuron v because all neurons can move around, not just the BMU. The second difference is the addition of the neighborhood function $\theta(u, v, s)$. This is a factor between zero and one. There are different ways to define this neighborhood. For example, the neurons could be ordered in a rectangular grid. The BMU and the eight surrounding points get a value of one, while the other points get zero. This means that every neuron will move together with its eight neighbors. This implies that two datapoints that have BMUs that are far apart in the grid are less similar than datapoints that are closer together in the grid. The neurons will effectively map the distribution of the data (Kohonen, 1982).

The neighborhood function follows a similar pattern as the learning rate. As the iterations continue, it starts to shrink, which causes neurons to be affected less by their neighbors. This allows for convergence of the map. There are many possible shapes of the neighborhood function. A neighborhood function can be Gaussian as well, or the grid can be hexagonal instead of rectangular.

5.4 Distance metric learning

Until now, we used Euclidean distance as a proxy for dissimilarity. But this is not the only way. It is possible to transform, weigh and combine different features to find a distance that separates classes as best as possible. Some features might be more predictive than others to separate classes, so it would make sense to weigh them more. The learning of this optimal transformation is called distance metric learning. There are countless techniques to learn a distance metric. For a general overview, we refer to Bellet et al. (2013). Here, we will only discuss one specific technique: distance metric learning through maximization of the Jeffrey divergence (DMLMJ). This technique was proposed by Nguyen et al. (2017) and yields Mahalanobis distance (Mahalanobis, 1936). Mahalanobis distance is given by:

$$d_M(x, y) = \sqrt{(x - y)^T S^{-1} (x - y)}, \quad (5.4)$$

where x and y are two vectors of a specific distribution and S is the covariance matrix. These vectors are not the original datapoints. With this method, we assume that the fea-

tures follow a multivariate Gaussian distribution that is different for each class. The distance can then be mapped as a standard deviation away from the centroid (a vector). As an example, consider a dataset with two dimensions: X and Y . We have a cluster $C1$ with a number of points and we want to know if a new point $P1$ belongs to this group. Then, we will check which points already belong to the cluster $C1$ and then we model the distances in the two dimensions as a multivariate normal distribution. Then, we can evaluate how likely it is that $P1$ follows this distribution. Graphically, this can be seen as ellipses of probability around the centroid of the cluster. The further out a point is, the less likely it is that it belongs to the cluster. If all points of the cluster are closer together, the ellipses grow closer together and a point far out becomes less likely to belong to the cluster. If the points grow farther apart, the ellipses also grow farther apart and a point that is far out is still likely to belong to the cluster, since we expect points to be this far out. If the features are correlated, the ellipses will start to rotate. This makes it possible that a point far away in either X or Y is unlikely to be in a cluster, but a point far away in both X and Y is more likely. This is shown visually in Figure 5.3.




Figure 5.3 Three different multivariate normal distributions. The more inwards the ellipse, the higher the chance a datapoint on this ellipse belongs to the cluster. The two figures on the left show negative correlation, a high value on one axis can be compensated by a low value on the other one. The right plot shows positive correlation. The right plot shows the circles further apart, which indicates the points show a higher variation around their centroid. Source: Andrew Ng, lecture notes CS229.

The previous two methods can be updated with distance metric learning. At the point where the distance to the centroid or neuron is calculated, the formula is updated to use the Mahalanobis distance instead of the Euclidean distance. Mahalanobis distance is usually saved in a weighing matrix. If this weighing matrix is the identity matrix, the Mahalanobis distance becomes equal to the Euclidean distance. Important to note is that DML is a supervised method, i.e. a labeled dataset will be needed for this method. Therefore, DML is only an expansion to the previous method of self-organizing maps. If there is no labeled dataset available, this expansion is not possible and we will have to revert to an Euclidean distance instead. When a DML is used, the method becomes semi-supervised.

5.5 Data transformation

All of the methods discussed before require the data to be numerical features where the concepts of distance have a clear meaning. This means sound data will need to be described with vectors of a fixed size. In this section two most common ways to do this will be discussed: spectrograms and scaleograms.

5.5.1 Spectrograms and Fourier transformation

A common way to visualize sounds is to represent them as a spectrogram. This is a graph that shows the frequencies in function of time. An example is given in Figure 5.4.




Figure 5.4 Example of a spectrogram. Species: *pipestrellus pipestrellus*. Data provided by INBO.

Spectrograms are usually created using a Fourier transformation. This is a transformation that decomposes a wave into the frequencies it is made up from. The method assumes that every sound is a combination of different sinusoidal curves with different frequencies. It has two parts: a real part and an imaginary part. The real part expresses how much a certain frequency is present in a sound and the complex argument describes the phase offset. Phase offset is the shift of a wave. A regular sinewave starts at zero amplitude, goes up to a maximum, then goes down to a minimum and then goes back to zero. If a wave has the same shape, but first goes down and then up, it has a 180 degree phase offset (Fourier, 1822). This is illustrated in Figure 5.5.

Fourier transformations can cause spectral leakage. This phenomenon is caused by the finiteness of both signal and frequency. A Fourier transformation will test out a finite number of sine waves. As long as the signal itself is an exact integer of periods long, this will work. But if this is not the case, sharp transitions are created and this causes the spectrum to smear out. Frequencies that are not in the signal will start to show up. Energy from one frequency leaks into the others, hence the name spectral leakage. Explained simply, if a signal contains exactly five waves, there is no problem. If a signal contains 5.2 waves, the Fourier transformation will try to fit in the extra one fifth of a wave somewhere, which causes frequencies to show up that are not actually there. The method does not know




Figure 5.5 Illustration of a phase shift. The phase shift is noted as θ and is 30 degrees in this case.

beforehand there are 5.2 waves because determining the number of waves is part of what the Fourier transformation does. Figure 5.5 illustrates the finiteness. Both signals have one full wave followed by a small portion of a second wave. This portion is about 30 degrees or one twelfth of a full wave. So the signal has roughly 1.08 waves, which makes it difficult to decompose properly. The use of windowing can lower the influence of spectral leakage. Under windowing, the amplitude along the edges of every waveform is artificially lowered, which prevents sharp transitions. There are different window functions. The most widely used ones are Hanning and Hamming windows (Harris, 1978).

5.5.2 Scaleograms: wavelet transformation

A special variant of this principle is the wavelet transform. A wavelet is a function in two variables. Similar to a sinewave, the amplitude is defined at every point in time. However, contrary to sinewaves, which can only differ in frequency, wavelets can also differ in shape and symmetry. This means that a wavelet can be crafted for a specific application. If music is analyzed, a wavelet can be crafted to respond to a specific note. Wavelets are combined using convolution (a mathematical operation used to combine different functions into one). A wavelet transformation produces a scaleogram instead of a spectrogram. Scaleograms have coefficients that state which wavelets need to be combined at which time. In a spectrogram, only the frequency is given, because all components are standard sinus functions (Addison, 2002). An illustration of a wavelet is given in Figure 5.6. In general, wavelets require more knowledge about the data involved.

Both wavelets and Fourier transforms are limited by the uncertainty principle. This puts a limit on the time and frequency resolution. To increase the time resolution, one must decrease the frequency resolution. If the resolutions are expressed as standard deviations, then their product cannot be smaller than $\pi/4$ (the Gabor limit, Benedicks (1985)). Standard deviation is a different way to express resolution, since both refer to uncertainty on




Figure 5.6 The Meyer wavelet (Meyer, 1990).

a specific value. Fourier transforms have both resolutions at a fixed value. However, a wavelet will increase time resolution as the frequency increases. This tends to work better for certain systems because actual hearing of animals works logarithmically. The difference between 50 and 51 Hz is easier to detect than the difference between 50 000 Hz and 50 001 Hz. This is not just for hearing, every sense works this way. Seeing the difference between one lion and two lions is very simple. Seeing the difference between a 100 and 101 lions is incredibly difficult even though we are talking about one extra lion in both cases. Due to this principle of nature, we can assume that bats will not be able to hear the difference between high frequencies as well as low frequencies. Hence, it would make sense to decrease the frequency resolution for higher frequencies.

5.6 Unbalanced data

When the distribution of the datapoints over the different classes is not uniform, a dataset is called unbalanced. This means that certain classes have a higher a priori probability of occurring. In the context of animal monitoring, this would mean that certain species are much more common than others. This comes with a few challenges, which we will discuss briefly. Our discussion is based on López et al. (2013).

The main problem with unbalanced data is that it can create bias towards majority classes. If 95% of the data belongs to species A and the rest belongs to twenty other species, a model could predict A for every datapoint and have an accuracy of 95%. This accuracy is not wrong, but it gives the wrong idea, because the mistakes all happen in the same area. Class A has a misclassification rate of zero, since every A is classified as A, while the other

classes have a misclassification rate of 100%. When the misclassification rates over the classes are not equal, the model is called biased towards one or more classes. This can be especially problematic if the small classes are the most interesting cases. In the context of conservation of species, it makes sense that the species that need conservation would be much rarer than other species. So the accuracy needs to be high in this area and not in the area of the common species. Medical research generally has this problem as well. The people who have a disease are usually a very small class, but classifying someone with a disease as healthy has disastrous consequences. In the reverse case, a second test can usually fix the mistake. In general, the rare datapoints carry more information than the common ones. So it will be important to focus more on these cases.

In some cases, different performance for different classes can arise due to lack of clear boundaries on a minority class. This is illustrated in Figure 5.7. This data has two features: horizontal and vertical axis and two classes: purple stars and blue dots. For both classes, we can draw a rectangle that contains all datapoints of one specific class. These are the class boundaries. It is clear that when extra data is added, the class boundaries of the minority class change dramatically, while those of the majority class change only slightly. In other words: the full range of variability is not captured for the minority class. If a new datapoint comes in that lies outside the initial narrow class boundaries, the model is very unlikely to classify it correctly. This will lower the performance for the minority class.




Figure 5.7 An illustration of the boundary problem in unbalanced data. The minority class is represented by purple stars and the majority class by blue dots. The left side shows 10% of the dataset and the right side shows the full dataset. Horizontal and vertical axis represent the two features of the dataset. The red rectangle represents the minority class boundaries, the black rectangle the majority class boundaries. Adapted from López et al. (2013).

These problems are inherent to all datasets, regardless of whether they are explicitly labeled or not. However, all potential solutions require data to be labeled. If the distribution of the classes is not known, it is not possible to take that distribution into account. Therefore, this subject will not be discussed any further. But it is an important detail to keep in mind when evaluating a model. If a bat species is more common, a higher performance is expected.

CHAPTER 6

PRE-PROCESSING AND CLUSTERING

In this chapter, the tool used to analyze the data is described. First the data used is described, followed by a discussion of pre-processing methods. Then, the analysis is discussed, which consists of two steps: exploration and clustering. A flowchart of all of these steps is given in Figure 6.1.

6.1 Data description

The data used in the case-study is provided by INBO (Research Institute for Nature and Forest). Most of these data come from domain Huizingen (a domain southeast of Brussels, Belgium). The data are collected using three different devices (IP-66, D-240X, RPA-03) over six different days ranging from May to September 2017. A portion of the data are labeled at the level of an individual audio file. These files have a size of a few megabytes, which corresponds to a few seconds of audio. A file may contain bats of different species. From these recordings, individual pulses are extracted manually. The field guide of Dietz (Dietz and Kiefer, 2017) is used as a guide to label individual pulses. There are a total of 634 audio files in this Huizingen dataset with 150 files unlabeled, 257 files labeled with certainty about the label and 227 files labeled with uncertainty about the label. Certain or confirmed means an independent expert verified the label. Unconfirmed means this is the label the operator gave, without any verification by a third party. Apart from this dataset, there are also data recorded along the shores of the Oevel channel (Antwerp). These data are recorded with the SM4BAT-FS detector (Wildlife Acoustics). This data is not labeled. Table 6.1 shows how common each species is in the labeled portion of the Huizingen dataset (both confirmed and unconfirmed labels).

From the table, it is clear that 94% of the recordings with one bat species are from a *Pipistrellus pipistrellus*. Other species only have a handful of recordings available and most species are missing entirely. As mentioned in Section 2.2.2, there are eighteen species that

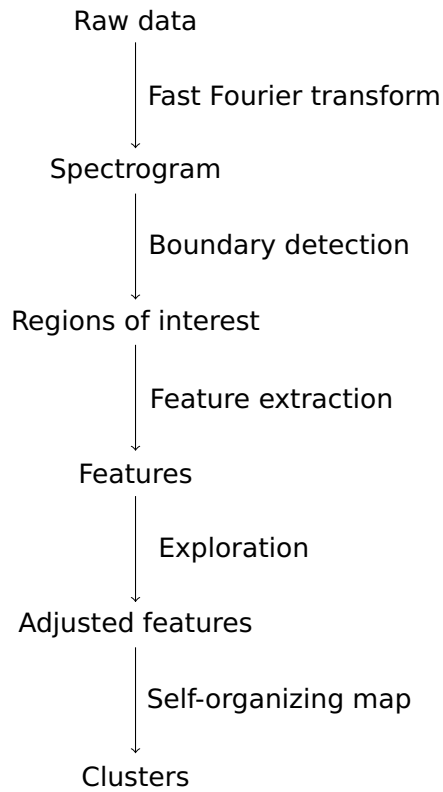


Figure 6.1 Flowchart of the different steps in clustering audio data.

Species	Number
Multiple species	107
Noise	25
Unknown Myotis species	3
<i>Eptesicus serotinus</i> (eser)	3
<i>Myotis daubentonii</i> (mdau)	5
<i>Nyctalus leisleri</i> (nlei)	1
<i>Nyctalus noctula</i> (nnoc)	3
<i>Pipistrellus nathusii</i> (pnat)	8
<i>Pipistrellus pipistrellus</i> (ppip)	337

Table 6.1 Species in the labeled portion of the Huizingen dataset.

are reported to the EU, and only six of those are present in our dataset, five of which have less than ten recordings. Because of this, an unsupervised method is preferred.

6.2 Pre-processing

6.2.1 Transformation

First, the data will be transformed. To that end, a spectrogram is used. This spectrogram is created with a fast Fourier transformation with a 1024 Hamming window. The same window is used in several papers about bat calls such as Mac Aodha et al. (2017). It is also commonly used for eco-acoustics in general. An example can be seen in Ulloa et al. (2018). This window results in a spectrogram with a frequency resolution of 375 Hz and a time resolution of 0.33 ms. An example of such a spectrogram is given in Figure 6.2a.




Figure 6.2 Example of a spectrogram. Calls are from a *Pipistrellus pipistrellus*.

The intensity is transformed into grayscale values between 0 and 255 (8 bits). A sliding window is used with a width of 200 ms and a 50 ms overlap. This is necessary because otherwise a loud noise in a recording could drown out all other sounds in the same recording. Especially in longer recordings, this could be a problem. Using trial and error, 200 ms was found to be a suitable value. Next, a subtraction method is used. The mean of all grayscale values within a window is taken and this value is subtracted from every value. If an image is very noisy, this mean value will be large, and hence a lot of noise will be removed. If the image is not noisy, the mean is basically zero and this will make no difference. Figure 6.2b shows what would happen if subtraction was not used. Clearly, there is more background noise here.

6.2.2 Region of interest extraction

Next, regions of interest (ROIs) are extracted from the spectrograms. In order to do this, images are converted to a binary form. To that end, a threshold of 25 is used, or 10% of the maximum intensity. This value was determined by trial and error. Our previous subtraction method will be important here to clear away noise. After the image is converted to binary, the boundary detection algorithm from Suzuki and Abe (1985) is applied. This algorithm will detect ROIs in spectrograms. These ROIs will be restricted to have a minimum size of 5 kHz (4875 Hz or 13 pixels) and 1 ms (3 pixels). This will allow us to extract real signals from noise efficiently since the vast majority of noise will not reach a range of 5 kHz. These boundaries were determined empirically. An illustration of the boundary detection is shown in Figures 6.3a and 6.3b.

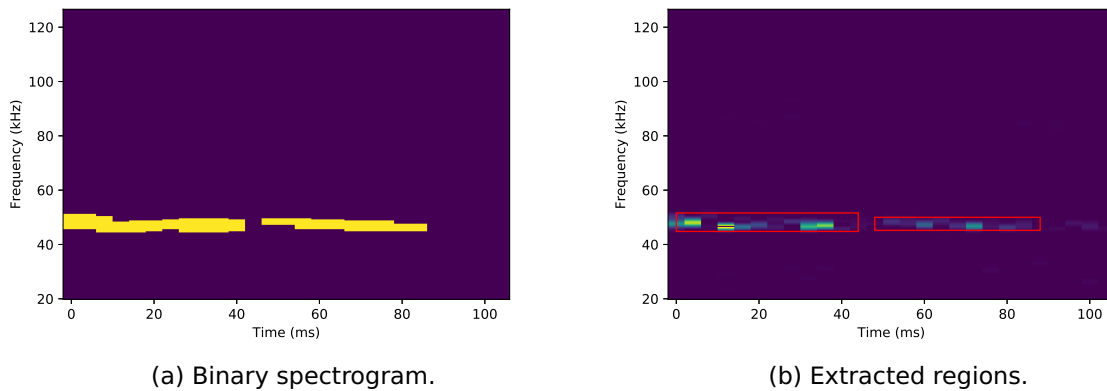


Figure 6.3 Example of a binary spectrogram. Calls are from a *Pipistrellus pipistrellus*.

Every region will now be rescaled individually to a range between 0 and 255. This ensures that there is no substantial difference between a loud call and a quieter call of the same bat.

6.2.3 Feature extraction

In the final pre-processing step, features will be extracted. In this context, a feature is a number that describes an aspect of the region. There are two kinds of features: positional and shape features.

Positional

The x-coordinates and y-coordinates in a spectrogram relate to time and frequency. This means that the lowest y-value is also the lowest frequency in a certain bat call. This frequency will be informative about the species that made the call. Apart from minimum frequency, we can also compute: maximum frequency, average frequency, frequency range

and duration. Lastly, there is the peak frequency. The peak is the loudest part of the region. The relative timing and frequency of this peak can be informative about the bat species as well. This is illustrated in Figure 6.4.

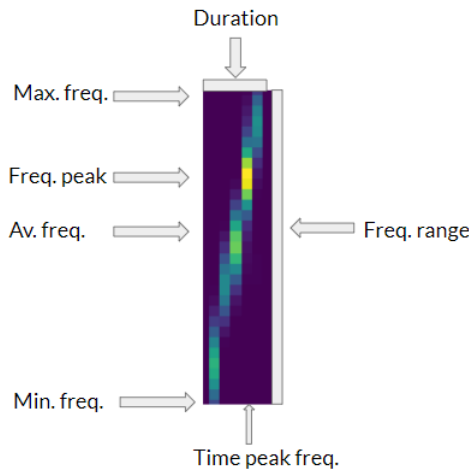


Figure 6.4 Illustration of the different frequency features that can be extracted from a region of interest.

Shape information

Shape information is always relative to something. So images will have to be compared with each other. The simplest metric to compare images is called the Mean Squared Error (MSE).

If both images have the same size, then every pixel has a matching pixel in the other image. The pixel at position $(3, 2)$ in image A will match with the pixel at position $(3, 2)$ in image B . We can now take the grayscale values of these pixels and compare them. For every pixel, we take the squared difference of the grayscale values. Then we take the average of all these values. This is called the MSE. For grayscale images, this will be a number between 0 and 255, but in general this number will be divided by the maximum pixel value to result in a standard value between zero and one. For color images, the MSE can be calculated separately for each color channel and averaged. The higher the MSE, the less similar two images are. For two images A and B , this is given by the following equation:

$$\text{MSE}(A, B) = \sum_{i=1}^n \frac{(\mu_{Ai} - \mu_{Bi})^2}{n}, \quad (6.1)$$

where μ_{Ai} is the value of pixel i in image A and n is the total number of pixels (Dosselmann and Yang, 2011). This method can be expanded further. Instead of one pixel at a time, a neighborhood of pixels can be considered. If this neighborhood is a simple square, then a neighborhood with radius 1 would be a central pixel and the 8 pixels surrounding it. If the radius is 2, we have 24 pixels around our central pixel. The average of these pixels can

be taken and these averages can be compared. If the neighborhood becomes larger, small differences in individual pixels will have less impact, while small differences in a group of close pixels will have more impact. A larger neighborhood will look at the general picture and focus less on tiny details. This gives us the following expanded equation:

$$\text{MSE}(A, B) = \sum_{j=1}^m \frac{(\mu_{Aj} - \mu_{Bj})^2}{m}, \quad (6.2)$$

where μ_{Aj} is average value of the pixels in window j of image A and m is the number of different windows. The strongest criticism of this method is the response to distortions and luminance shifting (Dosselmann and Yang, 2011). Above, we claimed that a higher MSE implies less similarity between two images. But this is not entirely true. Distortions can result in images that look very similar to the human eye, but still have a high MSE. The other way around, some images have a low MSE even though they are clearly very different. This is because different situations can result in the same MSE. For example, shifting every pixel in the same direction by 10% will result in an MSE of 0.1. Shifting every pixel in a random direction by a random number that averages to 10% will give the same MSE. In the second case, it is clear that the perceived change will be larger even though the MSE is the same. This is illustrated in Figure 6.5 where various distortions are applied to an image. Some distortions only have a narrow effect on the perception while others have a very big effect. However, all images have the same MSE compared to the original.




Figure 6.5 MSE and SSIM of a picture with various distortions applied. (a) Reference image. (b) Mean contrast stretch. (c) Luminance shift. (d) Gaussian noise. (e) Impulsive noise. (f) JPEG compression. (g) Blurring. (h) Spatial scaling (Wang and Bovik, 2009).

To combat this problem, the structural similarity index or SSIM was proposed (Wang et al., 2004). SSIM uses two values: the MSE and the covariance. Covariance is a joint variability in two variables. If we take an area in one image that consists of a high pixel surrounded

by low pixels and we compare this with another area with a high pixel surrounded by low pixels, we will find a high covariance. Even if the absolute values of both regions are very different, the covariance will still be high. Covariance measures whether the relationship between a pixel and its surroundings is similar in two areas. A uniform shift in values will result in a different MSE, but will keep the covariance identical. The formula to calculate covariance is given by:

$$\text{cov}(A, B) = \sum_{j=1}^m \frac{\sigma_{ABj}}{m}, \quad (6.3)$$

where σ_{ABj} is the covariance between image A and B in window j and m is the number of windows. To get to the SSIM, these two values will be combined. However, this will require a transformation. Right now, both values have a different scale and direction. Higher covariance indicates a better match, while higher MSE indicates a worse match. As such, we will have to invert MSE. Both values will also need to be normalized so they can be compared properly. This results in the following equation:

$$\text{SSIM}(A, B) = \sum_{j=1}^m \frac{(2\mu_{Aj}\mu_{Bj})(2\sigma_{ABj})}{m(\mu_{Aj}^2 + \mu_{Bj}^2)(\sigma_{Aj}^2 + \sigma_{Bj}^2)}, \quad (6.4)$$

where σ_{Bj} is the pixel variance in window j of image A , σ_{ABj} is the covariance between image A and B in window j , μ_{Aj} is the average pixel value in window j of image A and m is the number of windows. This equation has two components. The MSE value lies between zero and one. One means that every window in both images shows an identical mean. Zero would indicate the mean could not be more different. One image would be completely black and the other image completely white. The covariance value lies between minus one and one. Minus one indicates an inverse correlation: one region shows a high pixel value surrounded by lower pixel value and the other region shows a low pixel value surrounded by higher pixel values. One indicates perfect correlation. Under perfect correlation, the second image can be constructed from the first image using the same operation on every window of pixels. If one image is black and the other one white, we would have perfect correlation. While they could not be more different in terms of absolute error, in terms of structure they are identical. This way, the SSIM balances structure and absolute error.

There is one detail that needs to be fixed. During the normalization, we divided by $\mu_{Aj}^2 + \mu_{Bj}^2$ and $\sigma_{Aj}^2 + \sigma_{Bj}^2$. In certain cases, these values could become zero. For two pure white images for example, both of these values would be zero. To prevent dividing by zero, a constant is usually added to both nominator and denominator. This results in the final equation:

$$SSIM(A, B) = \sum_{j=1}^m \frac{(2\mu_{Aj}\mu_{Bj} + c_1)(2\sigma_{ABj} + c_2)}{m(\mu_{Aj}^2 + \mu_{Bj}^2 + c_1)(\sigma_{Aj}^2 + \sigma_{Bj}^2 + c_2)}, \quad (6.5)$$

with c_1 and c_2 small arbitrary positive numbers (e.g. 0.01). This final metric is called the structural similarity index (SSIM). SSIM was developed by Wang et al. (2004) and is a widely used method in image processing. Compared to other methods like MSE and peak to noise signal ratio (PSNR), SSIM focuses more on perception and less on absolute errors. For more information on the mathematical properties of the SSIM, we refer to Brunet et al. (2012).

6.2.4 Reference library

To work with shape information, a reference library of bat calls is defined. These are regions that are extracted from audio files that were manually classified. Some examples of these references are given in Figure 6.6. Figures are saved using a 19 digit hash code. Folders contain the species code such as eser or ppip. We expect that when a new call comes in from the same bat, it will show a large SSIM with the reference image. This library can be expanded and adapted when necessary.

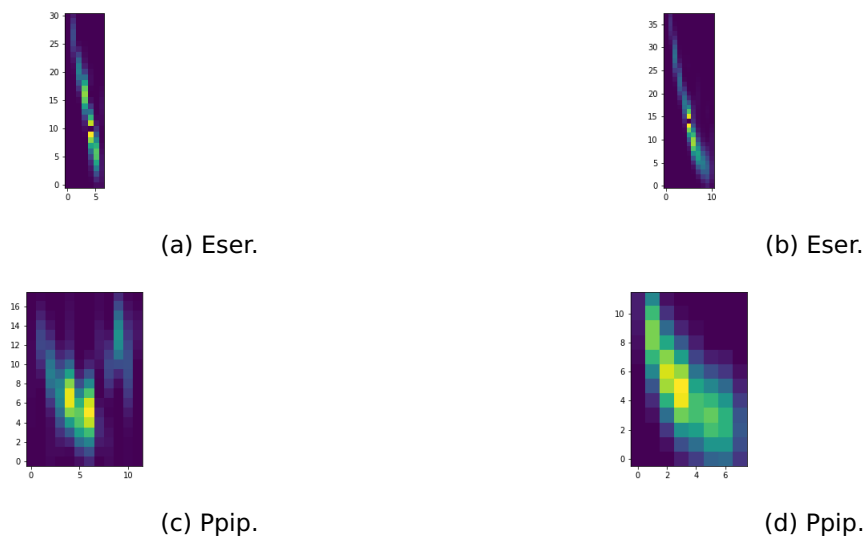


Figure 6.6 Examples of references.

6.3 Analysis

When all these steps are executed, the final result is that each region within a spectrogram is described by a list of numbers (features). These numbers contain information about the frequency and relative shape of the region. It is expected that similar bat calls will show similar features and thus these features can be used to classify sounds. As features,

we use seven frequency parameters plus the number of templates that are used from the reference library as shape factors. In the next sections, we will discuss exploration methods and classification methods. Exploration methods are methods to visualize and adapt the features. Classification methods connect a specific region to a species.

6.3.1 Exploration methods

Exploration methods are used to visualize the data. These can be used to test different features such as parameter settings, devices or locations. They can also be used to obtain an estimate of the expected performance of a classification method.

Multi-dimensional scaling


A multi-dimensional scaling (MDS) is a low-dimensional representation of high-dimensional data. The distance between different datapoints is calculated and approximated as good as possible in two dimensions. Using an MDS, we can represent how similar certain points are. The distance can be defined in different ways. In the case of a metric MDS (mMDS), the Euclidean distance is used. But if more information is known, certain important features can be weighed differently. Under the Euclidean distance metric, MDS is equivalent to principal coordinates analysis (PCoA) (Borg and Groenen, 2005). An example of an MDS is given in Figure 6.7a.

t-distribution stochastic neighboring embedding


t-distribution stochastic neighboring embedding (t-SNE) is analogous to an MDS, but uses a different principle (Maaten and Hinton, 2008). Instead of distance, this method fits a probability distribution. It can be useful to compare a TSNE to an MDS to make sure there is no influence of the visualization method. An example of a t-SNE is given in Figure 6.7b. From the figures, it is clear that the general picture is the same, but some details are different. In the example shown, TSNE trades smaller variability within a species (such as nlei) in favor of a larger distance between two big groups of species (ppip and pnat versus others). Note that neither method had access to these labels. The labels were only used to color the markers in the end.

6.3.2 Classification method: self-organizing map

As a classification method, a self-organizing map is used in combination with a DML matrix (Sections 5.3.2 and 5.4). Both the map and the DML-matrix will need to be fitted first. The



(a) MDS



(b) TSNE

Figure 6.7 Example of an MDS/t-SNE applied on six species of bats. Data provided by INBO.

DML matrix is fitted on labeled pulses. These labeled pulses can either come from a small dataset, or from the reference library. Because the references defined for the shape factors are also datapoints, they can be used to fit the DML. If this is done, a portion of these references cannot be used for shape factors anymore, because the number of datapoints used to fit the DML-matrix needs to outweigh the number of features. If this is not the case, the DML can start to overfit. To fit a DLM, the Jeffrey divergence is calculated (Nguyen et al., 2017).

The SOM is fitted on unlabeled data. These data can be the same data as the user wants to analyze, or come from a different but similar dataset. To find the BMU at each step, the Mahalanobis distance is used. This is given by the DML-matrix that was fitted before. Since a performance cannot be calculated with unlabeled data, a tuning set or cross-validation approach cannot be used to fit the parameters. Therefore, we will use trial and error to set the different parameters. Neurons are saved in a three dimensional matrix. The rows and columns are set by the user and determine how many neurons there are. The number of pages is the same as the number of features. The number of iterations n_{iter} is set to 10 000. The initial learning rate is set to 0.01 and decays using an exponential function with the exponent $\frac{-i}{n_{iter}}$ where i is the current iteration and n_{iter} are the total number of iterations. For the neighborhood, the radius r is used. Every neuron within a distance r^2 is considered part of the neighborhood. These neighboring neurons update their weight according to the influence factor, where influence is given by an exponential function with exponent $\frac{-D}{2r^2}$. D is the distance to the BMU. If D goes up, the influence factor goes down. If the radius is higher, the influence factors also become higher. Initially the radius is the highest dimension of the map divided by two. So a five by seven map would have 35 neurons and a radius of $7/2$. A six by six map would have 36 neurons and a radius of $6/2$, while a thirty-five by one map would have a very high radius of $35/2$. This allows the user to influence the radius by changing the configuration of the neurons. The radius decays exponentially with the exponent $\frac{-i}{t_{const}}$, where t_{const} is the time constant, given by the number of iterations divided by the log of the initial radius. If there are more iterations, the time constant becomes higher and the radius decays more slowly. If the initial radius is higher, the time constant decreases and the radius decays faster. Finally, there are the number of neurons. As mentioned before, number of neurons is defined as rows and columns of the matrix. More neurons will give a more detailed map, but will take much longer to fit.

6.3.3 Visualization tool

After a SOM is fitted, data can be analyzed using a visualization tool. This tool will group the datapoints together according to their BMU and order them according to the distance to this

neuron. An expert can now go through the pulses in each neuron and label them. Because of the clustering, it is likely that only a limited number of datapoints will need to be labeled manually to draw a conclusion for all datapoints matching with a neuron. Furthermore, if maps are reused, specific neurons can be tied to one or more species. Because datapoints are ordered according to distance, an outlier is likely to show a large distance to their BMU, which can make it easier to detect rare species. Neurons and datapoints can also be plotted in an MDS format, which makes it possible to visually inspect how well the neurons cluster the datapoints together. This is further illustrated in Chapter 7.

CHAPTER 7

RESULTS AND DISCUSSION

In this chapter, the previous methods will be evaluated to determine their performance. Due to the low amount of labeled data available, a formal evaluation is not possible. Therefore, two informal evaluations will be done. The first evaluation will walk through an analysis of a few real files that are unlabeled. The second evaluation will set up a few experiments to test several hypotheses. To test these hypotheses, a K -nearest neighbors method is used along with the Cohen's kappa.

7.1 Application of the tool

To apply the visualization tool, six unlabeled recordings are randomly selected. A standard reference library was defined before. This reference library contains five common species. We used 17 *Eptesicus serotinus* (eser), 6 *Myotis daubentonii* (mdau), 21 *Nyctalus leisleri* (nlei), 18 *Pipistrellus nathusii* (pnat) and 38 *Pipistrellus pipistrellus* (ppip) pulses. These are pulses, not entire recordings. Some pulses might come from the same recording, but to ensure a minimum variability, every species has pulses that come from at least two different locations, two different devices and two different days. There are also 10 eser and 11 ppip references extra to fit the DML (the DML matrix is fitted on the reference library). Then, a ten by ten map is fitted on the unlabeled recordings.

In total, there are 726 pulses over these six recordings. These pulses are divided over the 100 neurons. The neuron with the highest number of matches is neuron (9, 5) with 28 matches. The lowest number of matches is zero for several neurons.

Now, we can go through every neuron and see what is inside. Figure 7.1 shows datapoints 0, 5, 15 and 25 of neuron (9, 5). Remember that points are ordered according to distance to the neuron and the number shows their rank, 0 is the closest, 27 is the farthest away. From these points, it is clear that neuron (9, 5) contains noisy pulses of one or more bat species around 25-30 kHz. It is no longer necessary to study each of the 28 pulses individually. From a handful of pulses, the general story can be easily seen since all of these pulses are quite similar.




Figure 7.1 Points 0, 5, 15 and 25 from neuron (9, 5). This neuron contains noisy pulses of one or more bat species around 25-30 kHz.

This process can be repeated for every neuron. Neuron (2, 0) for example has 9 matches. Figure 7.2 shows matches 0, 2, 6 and 8. These pulses are cleaner. Given the shape and the frequency of the pulses, this is likely a *Pipestrellus* species, but this should be confirmed by an expert to be certain. Once an expert has labeled these pulses, this neuron is labeled too. If the same map is used, future matches with different data are likely to be the same species as before or a similar species.




Figure 7.2 Points 0, 2, 6 and 8 from neuron (2, 0). Given the shape and the frequency of the pulses, this is likely a *Pipestrellus* species, but this should be confirmed by an expert to be certain.

The distance between neurons can also be studied. If there is a different neuron that is very close to neuron (2, 0), this could contain the same species or a similar species. This principle can be taken further. On a sufficiently large map, regions of species or genera will start to emerge. We will not try to make a large map like this, because as mentioned before (Section 6.1), there is not enough good data of specific species available to do this.

Furthermore, this is something that should be done by an expert due to the complexity of bat pulses.

7.2 Evaluation experiments

7.2.1 Description of the datasets

For the evaluation experiments, two datasets will be used. The first dataset is provided by INBO and contains sounds from various locations in Flanders. This dataset is used to build the model. Templates are defined for shape features (Sections 6.2.3 and 6.2.4) and the Mahalanobis distance (Section 5.4). The second dataset is a completely independent dataset from Barataud (Barataud, 2015). This dataset contains bat sounds from Europe. The second dataset contains bats that are not present in Flanders or the same bats in very different environments. With this dataset, we can try to evaluate how the method responds to entirely new data.

7.2.2 Method: K -nearest neighbors and Kappa statistic

As an evaluation method, the K -nearest neighbors method is used. First, the features are calculated for every pulse. If we have the features, we can calculate the distance between datapoints. In this case, Mahalanobis distance is used. Then, for every point, the K -nearest neighbors are considered. So if K is set to 3 and we consider a datapoint belonging to species A, we will check if the three datapoints closest to this point (the nearest neighbors) belong to the same species. The more neighbors belong to the same species, the better the clustering is. If this is repeated for every point, a global score can be calculated for every species. If we do this, we may for example find a score of 66% for species A. This means that on average, two of the three nearest neighbors belong to the same species. This is called the *relative observed agreement*, indicated with $P(A)$.

This value has one problem, however. It does not say anything about the a priori chance. If there are 100 points belonging to species A and only 10 points belonging to species B, then there would be a very high chance that three neighbors belong to species A, even if the points are randomly distributed. So in this situation, two out of three neighbors for species A would actually be a very bad score. But for species B, two out of three would be a very good score. To capture this, the Cohen's kappa statistic can be used (Smeeton, 1985). The Cohen's kappa is computed as:

$$\kappa(A) = \frac{P(A) - P(E)}{1 - P(E)}. \quad (7.1)$$

In this equation $\kappa(A)$ is Cohen's kappa for species A , $P(A)$ is the relative observed agreement for species A and $P(E)$ is the expected agreement by chance. At the end, we divide by $1 - P(E)$. This sets the highest possible κ equal to one and makes it possible to compare different values in different settings. Note that κ is calculated per species. Multiple *kappa*-values will be needed to draw a conclusion about the clustering as a whole.

To understand κ further, the equation can be applied to the previous example. In our example, $P(E)$ would be $99/109$ for species A , since there are a total of 109 neighboring points and 99 of those belong to species A , therefore the chance that a random neighbor belongs to species A is $99/109$ or about 91%. κ now has a clear meaning. If more than 91 points out of every 100 neighbors belong to species A , κ is positive and the method does better than would be expected based on the distribution of the datapoints. There is a clustering effect. For a random point, the chance that it belongs to species A becomes higher if it has neighbors that belong to species A . If exactly 91 neighbors belong to class A , κ is zero. This is the value that would be expected if all points were distributed randomly. If there are less than 91 neighbors, κ becomes negative. This value would indicate a 'reverse clustering effect'. The fact that a datapoint is a neighbor of a different datapoint of species A now *decreases* the chance this point also belongs to species A .

7.2.3 Self-organizing maps

Self-organizing maps are not used in the evaluation experiments because the performance will be different depending on the size. If there are more neurons available, a better clustering is possible. Furthermore, self-organizing maps are an unsupervised method. Some species might be represented by more neurons than others, which would give them a better performance. If ten datapoints of species A and ten datapoints of species B are clustered by three neurons, there is a reasonable chance some neurons will contain datapoints of both species A and B . If twenty neurons are used, every datapoint can have their own neuron resulting in no neurons with multiple species. Because of this, self-organizing maps are more difficult to evaluate. Instead, the K -nearest neighbors method is used. But both of these methods are related. High κ s would indicate a good clustering and therefore, a good self-organizing map.

Dataset	Species	Number	P	Cohen's kappa
Training	eser	12	0.58	0.44
Training	nlei	10	0.7	0.62
Training	ppip	24	0.97	0.94
Validation	eser	12	0.61	0.47
Validation	nlei	10	0.73	0.66
Validation	ppip	24	1.0	1.0

Table 7.1 Results of experiment 1. Top half shows the training dataset, bottom half shows the validation dataset. The bat species are: *Eptesicus serotinus* (eser), *Nyctalus leisleri* (nlei) and *Pipistrellus pipistrellus* (ppip). P is the relative observed agreement. K is selected as 3 in this experiment.

7.3 Experiments

7.3.1 Experiment one: generalization within species

In the first experiment three species are evaluated. For the species *Eptesicus serotinus* (eser), *Nyctalus leisleri* (nlei) and *Pipistrellus pipistrellus* (ppip), three common species in Flanders, pulses are extracted from the INBO dataset. In total, we have 24 eser pulses, 20 nlei pulses and 48 ppip pulses. Variability is somewhat limited due to the low number of labeled audio data available, but every species has pulses from at least two different locations, two different devices and two different days. These pulses are then split in two sets. The first set is the training set, used to compute the features (Section 6.2.3) and fit the DML-matrix (Section 5.4). The second set is the validation set, which is not used for any fitting. The training set is split again. The first half is used for the features. In total, we have 30 features, 7 standard frequency factors and 23 shape factors from the 23 templates in the first half of the training dataset (6 eser, 5 nlei and 12 ppip). The second half of the training dataset is used along with the first half to fit a DML-matrix. This ensures that the number of datapoints to fit the DML-matrix (46) outweighs the number of features (30), which is very important. If there are only slightly more datapoints, the DML-matrix tends to overfit on those datapoints.

Now, the K -nearest neighbors method is used with three neighbors. First on the training dataset then on the independent validation dataset. The results are summarized in Table 7.1.

As is clear from the table, the performance does not drop from the training dataset to the validation dataset. In fact, it even goes up slightly for all three species. We can conclude that it is possible to extrapolate between different recordings in Flanders. The features and DML matrix defined in one recording are still valid for different recordings on other days in other locations within Flanders. The performance on all three species is also quite large. The model predicts the label with a higher accuracy than would be expected just based on

how many datapoints each species has. The ppip species shows maximum performance on the validation dataset and almost maximum performance on the training dataset. This can be explained rather easily. The average frequency of the ppip bat lies around 45 kHz while both the nlei and the eser lie around 25 kHz. Therefore, the pulses of the ppip can be easily separated on frequency alone. Nlei and eser bats are quite similar in both shape of the pulse and frequency, so these will be more difficult to separate. However, the κ s are still positive.

7.3.2 Experiment two: influence of K and DML

For the second experiment, the same data is used as in experiment 1. But now, the influence of certain parameters will be tested. Firstly, the numbers of neighbors K considered in the KNN. For the data we used, K can vary between 1 and 9, since nlei has 10 datapoints and can therefore only have 9 neighbors of the same class. Figure 7.3 shows Cohen's kappas for all values of K . From this figure, it is clear the influence of K varies. For the training eser datapoints there does not seem to be a large influence, but the evaluation eser dataset shows a clear maximum performance around 3 and 4. What is happening here is that there are two points of the eser cluster that lie very close to the nlei cluster. So close that their first few neighbors are nlei datapoints. Therefore, with only 1 or 2 neighbors considered, these points have a very low score and pull down kappa substantially. More distant neighbors are eser points again, so if more neighbors are considered, the influence of these initial outliers is reduced. Near the end, kappa goes slightly down again because the 7th, 8th and 9th datapoints are less likely to be correct given there are only 9 datapoints of the same class to start with. To reduce these two edge effects, it is best to pick a K somewhere around the middle. For future experiments, a K of 5 will be chosen.

Secondly, the influence of the DML-matrix will be evaluated. The validation dataset will be analyzed twice. Once with the DML-matrix learned from the training dataset and a second time using the identity matrix, which results in the usual Euclidean distance. A K of 5 is used. The results of this experiment are shown in Table 7.2.

The table shows no effect for ppip and eser, but a substantial drop in performance for nlei. As mentioned before, in comparison, ppip is relatively easy to classify because the frequency is substantially different from the other two bats. Therefore, a DML-matrix will not be needed to separate these points. But for nlei, there is higher similarity and without a DML-matrix, they become much more difficult to cluster together.

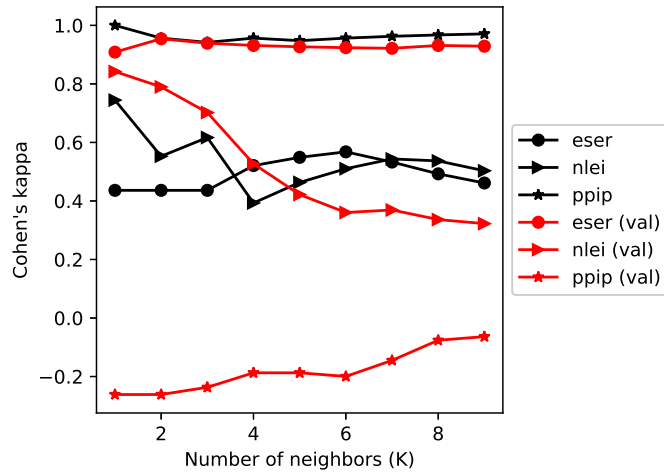


Figure 7.3 Influence of K on Cohen's kappa. Val refers to the validation dataset. The bat species are: *Eptesicus serotinus* (eser), *Nyctalus leisleri* (nlei) and *Pipistrellus pipistrellus* (ppip).

Distance	Species	Number	P	Cohen's kappa
Mahalanobis	eser	12	0.68	0.57
Mahalanobis	nlei	10	0.68	0.59
Mahalanobis	ppip	12	1.0	1.0
Euclidean	eser	12	0.68	0.57
Euclidean	nlei	10	0.48	0.34
Euclidean	ppip	12	1.0	1.0

Table 7.2 Results of experiment 2. Top half shows the validation dataset using Mahalanobis distance between the features, bottom half shows the validation dataset using Euclidean distance. The bat species are: *Eptesicus serotinus* (eser), *Nyctalus leisleri* (nlei) and *Pipistrellus pipistrellus* (ppip). P is the relative observed agreement. The number of neighbors (K) is five.

Dataset	Species	Number	P	Cohen's kappa
Training	eser	12	0.67	0.55
Training	nlei	10	0.58	0.46
Training	ppip	12	0.98	0.95
Validation	hsav	13	0.94	0.93
Validation	msch	52	0.73	0.42
Validation	ppyg	17	0.06	-0.19

Table 7.3 Results of experiment 3. Top half shows the training dataset (INBO), bottom half shows the validation dataset (Barataud). The bat species are: *Eptesicus serotinus* (eser), *Nyctalus leisleri* (nlei), *Pipistrellus pipistrellus* (ppip) *Hypsugo savii* (hsav), *Myotis schreibersii* (msch) and *Pipistrellus pygmaeus* (ppyg). P is the relative observed agreement. The number of neighbours (K) is five.

7.3.3 Experiment three: generalization to different species

For the final experiment, we will test whether the model can recognize new species it has never seen before. The Barataud dataset is used (Barataud, 2015). From this dataset, pulses are extracted for three common bats: *Hypsugo savii* (hsav), *Myotis schreibersii* (msch) and *Pipistrellus pygmaeus* (ppyg). The features and DML-matrix are computed using the original dataset from experiment one. Then, a KNN is calculated using five neighbors. The results of this experiment are summarized in Table 7.3. As is clear from the table, hsav shows a very high performance, msch shows a moderate performance and ppyg shows a very bad performance. So even though the features and DML-matrix were computed with different bat species, it is still possible to cluster new bat species to some extent.

CHAPTER 8

CONCLUSIONS AND FUTURE

PERSPECTIVES

The goal of this research was to create a tool that can cluster bat pulses together to allow for easier labeling. Due to the lack of labeled data, a formal evaluation cannot be done. However, the ad hoc application of the visualization tool showed promising results. Clearly, the pulses that are clustered together are very similar. The evaluation experiments confirmed this too. Within the same species, very high performances were found. For new species, the results were mixed. Cohen's kappas were 0.93 and 0.42 for *hsav* and *msch*, but *ppyg* showed a negative kappa. The large differences between these species indicate that care must be taken when generalizing to completely new species. There were only three species tested, so it is not clear how widely these principles can be applied to other species. Further testing with labeled data could highlight which species are easy and which ones are harder to classify. While we should not generalize too much from this rather limited experiment, it does indicate that clustering of pulses in different species is possible. Furthermore, this can also work if no examples are given from a certain species.

The use of SOM, DML and reference libraries allows for flexibility in the method. Depending on the application, the SOM can be expanded or reduced and the number of references can be adapted. It is also possible to use a multi-stage SOM. In this case, several neurons could be labeled as a group of closely related species. These pulses could then be clustered further with a second SOM. Due to the lack of labeled data, this was not done in this research. However, because there is still manual labeling involved with this tool, more and more labeled data is created during use. This data can then be used to further improve the model. Perhaps in the future, a fully supervised tool could be created that no longer requires expert input, or at least very minimal input.

Further pre-processing could also improve the model. A simple tree-based method could discard pulses that are likely to be noise. This is not done right now, because there is no sufficiently large labeled dataset of noise and bats available, so the risk of discarding a rare bat is too large. But if more labeled data is collected, this could become possible in the future. The use of specific wavelets that optimally separate certain bat species could also

improve the model. Lastly, metadata can be added to the model. Right now, the model uses individual pulses, but far more data can be extracted from bat detectors. The exact time related to sundown can be interesting, because some bats come out earlier in the night than other bats. There may be seasonal effects as well. Weather data could also influence which bats come out during a night. Furthermore, the timing between different pulses of the same individual can be useful to take into account for identification. Lastly, information about the habitat can be useful as well. Right now, data is analyzed blindly, but we know that some species are far more likely to be present in specific habitats.

BIBLIOGRAPHY

- Adams, A. M., Jantzen, M. K., Hamilton, R. M., and Fenton, M. B. (2012). Do you hear what I hear? Implications of detector selection for acoustic monitoring of bats. *Methods in Ecology and Evolution*, 3(6):992–998.
- Addison, P. S. (2002). *The Illustrated Wavelet Transform Handbook*. Institute of Physics, 1st edition.
- Altringham, J. (2011). *Bats: From Evolution to Conservation*. Oxford University Press, USA, 1st edition.
- Ancillotto, L., Serangeli, M. T., and Russo, D. (2013). Curiosity killed the bat: Domestic cats as bat predators. *Mammalian Biology*, 78(5):369–373.
- Baerwald, E. F., D’Amours, G. H., Klug, B. J., and Barclay, R. M. (2008). Barotrauma is a significant cause of bat fatalities at wind turbines. *Current Biology*, 18(16):695–696.
- Barataud, M. (2015). *Acoustic ecology of European bats. Species Identification and Studies of Their Habitats and Foraging Behaviour*. Biotope, 1st edition.
- Barclay, R. M. R. (1999). Bats are not birds: A cautionary note on using echolocation calls to identify bats: A comment. *Journal of Mammalogy*, 80(1):290–296.
- Barlow, K., Briggs, P., Haysom, K., Hutson, A., Lechiara, N., Racey, P., Walsh, A., and Langton, S. (2015). Citizen science reveals trends in bat populations: The national bat monitoring programme in Great Britain. *Biological Conservation*, 182:14 – 26.
- Bellet, A., Habrard, A., and Sebban, M. (2013). A survey on metric learning for feature vectors and structured data. *CoRR*, preprint arXiv:1306.6709.
- Benedicks, M. (1985). On Fourier transforms of functions supported on sets of finite Lebesgue measure. *Journal of Mathematical Analysis and Applications*, 106(1):180–183.
- Boonman, A., Bumrungski, S., and Yovel, Y. (2014). Nonecholocating fruit bats produce biosonar clicks with their wings. *Current Biology*, 24:2962–2967.
- Boonman, A. M., Parsons, S., and Obrist, M. K. (2000). Advantages and disadvantages of techniques for transforming and analyzing Chiropteran echolocation calls. *Journal of Mammalogy*, 81(4):927–938.

- Borg, I. and Groenen, P. (2005). *Modern multidimensional scaling: theory and applications*. New York: Springer-Verlag, 2nd edition.
- Boyles, J. G., Cryan, P. M., McCracken, G. F., and Kunz, T. H. (2011). Economic importance of bats in agriculture. *Science*, 332(6025):41–42.
- Brigham, R., Kalko, E., Jones, G., Parsons, S., and Limpens, H. (2004). *Bat Echolocation Research. Tools, Techniques, and Analysis*. Bat Conservation International.
- Brunet, D., Vrscay, E. R., and Wang, Z. (2012). On the mathematical properties of the structural similarity index. *IEEE Transactions On Image Processing*, 21:1488–1499.
- Chapelle, O., Schölkopf, B., and Zien, A. (2006). *Semi-supervised learning*. MIT Press, 1st edition.
- Chesmore, E. (2001). Application of time domain signal coding and artificial neural networks to passive acoustical identification of animals. *Applied Acoustics*, 62(12):1359–1374.
- Chiu, C. and Moss, C. (2005). The role of tragus on echolocating bat, *Eptesicus fuscus*. *The Journal of the Acoustical Society of America*, 117(4):2468–2468.
- Chollet, F. (2017). *Deep Learning with Python*. Manning Publications Co., Greenwich, CT, USA, 1st edition.
- Corcoran, A. J., Barber, J. R., and Conner, W. E. (2009). Tiger moth jams bat sonar. *Science*, 325(5938):325–327.
- Corcoran, A. J. and Moss, C. F. (2017). Echolocating bats cry out loud to detect their prey. *Journal of Experimental Biology*, 220:4554–4566.
- Council of European Union (1992). Council Directive 92/43/EEC. http://ec.europa.eu/environment/nature/legislation/habitatsdirective/index_en.htm, accessed: 2019-03-07.
- Cutnell, J. D. and Johnson, K. W. (1997). *Physics*. John Wiley & Sons Inc, 4rd edition.
- Dauwe, S., Oldoni, D., De Baets, B., Van Renterghem, T., Botteldooren, D., and Dhoedt, B. (2014). Multi-criteria anomaly detection in urban noise sensor networks. *Environmental Science: Processes Impacts*, 16:2249–2258.
- Davies, M. (2016). Bat bill put forward to protect churches. *Church Times*.
- de Oliveira, H. F. M., de Camargo, N. F., Gager, Y., and Aguiar, L. M. S. (2017). The response of bats *mammalia : Chiroptera* to habitat modification in a neotropical savannah. *Tropical Conservation Science*, 10:1–14.
- Denzinger, A. and Schnitzler, H. (2013). Bat guilds, a concept to classify the highly diverse foraging and echolocation behaviors of microchiropteran bats. *Front Physiology*, 4:164.

BIBLIOGRAPHY

- Dietz, C. and Kiefer, A. (2017). *Veldgids: Vleermuizen van Europa (Die Fledermause Europas)*. KNNV uitgeverij, Zeist.
- Dosselmann, R. and Yang, X. D. (2011). A comprehensive assessment of the structural similarity index. *Signal, Image and Video Processing*, 5(1):81–91.
- Drake, F. (2015). Bill reignites dispute over bats in churches. *Church Times*.
- Eionet, European Environment Agency (2013). Biogeographical regions, Natura 2000. https://biodiversity.eionet.europa.eu/activities/Natura_2000/chapter1, accessed: 2019-03-23.
- Eionet, European Environment Agency (2014). Article 17, 2012 reports. <http://art17.eionet.europa.eu/article17/reports2012/>, accessed: 2019-03-03.
- Eiting, T. P. and Gunnell, G. F. (2009). Global completeness of the bat fossil record. *Journal of Mammalian Evolution*, 16(3):151–173.
- Estivill-Castro, V. (2002). Why so many clustering algorithms: A position paper. *SIGKDD Explorations Newsletter*, 4(1):65–75.
- Ey, E. and Fischer, J. (2009). The "acoustic adaptation hypothesis" - a review of the evidence from birds, anurans and mammals. *Bioacoustics*, 19:21–48.
- Farina, A., Pieretti, N., and Piccioli, L. (2011). The soundscape methodology for long-term bird monitoring: A Mediterranean Europe case-study. *Ecological Informatics*, 6(6):354 – 363.
- Fenton, M. B. (2012). Bats and white-nose syndrome. *Proceedings of the National Academy of Sciences*, 109(18):6794–6795.
- Florida Museum of Natural History, University of Florida (2018). Artificial bat house. <https://www.floridamuseum.ufl.edu/bats/>, accessed: 2019-03-03.
- Fourier, J.-B.-J. (1822). *Théorie analytique de la chaleur*. Paris: F. Didot.
- Goodman, S. M. (2006). Hunting of Microchiroptera in south-western Madagascar. *Oryx*, 40(2):225–228.
- Groffen, J., Rethus, G., and Pettigrew, J. (2016). Promiscuous pollination of Australia's baobab, the boab, *Adansonia gregorii*. *Australian Journal of Botany*, 64:678–686.
- Harris, F. J. (1978). On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66(1):51–83.
- Hartigan, J. A. and Wong, M. A. (1979). Algorithm AS 136: A K-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108.

- Hodgkison, R., Balding, S. T., Zubaid, A., and Kunz, T. H. (2006). Fruit bats (chiroptera: Pteropodidae) as seed dispersers and pollinators in a lowland Malaysian rain forests. *Biotropica*, 35(4):491–502.
- Holland, R. A., Waters, D. A., and Rayner, J. M. V. (2004). Echolocation signal structure in the megachiropteran bat *Rousettus aegyptiacus* Geoffroy 1810. *Journal of Experimental Biology*, 207(25):4361–4369.
- Hoy, R. R. and Fay, R. R. (1998). *Comparative Hearing: Insects*. Springer-Verlag New York, 1st edition.
- Jones, G. (2009). Echolocation. *Current Biology*, 15:484–488.
- Jones, G. and Holderied, M. W. (2007). Bat echolocation calls: adaptation and convergent evolution. *Proceedings of the Royal Society of London B: Biological Sciences*, 274(1612):905–912.
- Jones, G., Jacobs, D., Kunz, T., Willig, M., and PA, R. (2009). Carpe noctem: the importance of bats as bioindicators. *Endangered Species Research*, 8:93–115.
- Kasso, M. and Balakrishnan, M. (2013). Ecological and economic importance of bats (order chiroptera). *ISRN Biodiversity*, 2013:9.
- Kim, S. and Stephen, A. (2018). Extreme heat wipes out almost one third of Australia's spectacled flying fox population. *ABC news*.
- Kober, R. and Schnitzler, H. (1990). Information in sonar echoes of fluttering insects available for echolocating bats. *The Journal of the Acoustical Society of America*, 87(2):882–896.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43(1):59–69.
- Krause, B. (1993). The niche hypothesis: A virtual symphony of animal sounds, the origins of musical expression and the health of habitats. *Soundscape Newsletter*, 6:6–10.
- Krause, B. and Farina, A. (2016). Using ecoacoustic methods to survey the impacts of climate change on biodiversity. *Biological Conservation*, 195:245–254.
- Lawrence, B. D. and Simmons, J. A. (1982). Measurements of atmospheric attenuation at ultrasonic frequencies and the significance for echolocation by bats. *The Journal of the Acoustical Society of America*, 71(3):585–590.
- Lizarro, D., Isabel Galarza, M., and Aguirre, L. (2010). Tráfico y comercio de murciélagos en Bolivia. *Revista Boliviana de Ecología y Conservación Ambiental*, 27:63–75.
- Lloyd, S. (1982). Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137.

BIBLIOGRAPHY

- López, V., Fernández, A., García, S., Palade, V., and Herrera, F. (2013). An insight into classification with imbalanced data: Empirical results and current trends on using data intrinsic characteristics. *Information Sciences*, 250:113–141.
- Maaten, L. v. d. and Hinton, G. (2008). Visualizing data using t-SNE. *Journal of machine learning research*, 9(Nov):2579–2605.
- Mac Aodha, O., Gibb, R., Barlow, K., Browning, E., Firman, M., Freeman, R., Harder, B., Kinsey, L., Mead, G., Newson, S., Pandourski, I., Parsons, S., Russ, J., Szodoray-Paradi, A., Szodoray-Paradi, F., Tilova, E., Girolami, M., Brostow, G., and Jones, K. E. (2017). Bat detective - deep learning tools for bat acoustic signal detection. *bioRxiv*, 14:1–19.
- MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, pages 281–297, Berkeley, Calif. University of California Press.
- Mahalanobis, P. C. (1936). On the generalised distance in statistics. *Proceedings of the National Institute of Sciences of India*, 2:49–55.
- Matteson, K. C., Grace, J. B., and Minor, E. S. (2013). Direct and indirect effects of land use on floral resources and flower-visiting insects across an urban landscape. *Oikos*, 122(5):682–694.
- McCauley, R. D., Gavrilov, A. N., Jolliffe, C. D., Ward, R., and Gill, P. C. (2018). Pygmy blue and Antarctic blue whale presence, distribution and population parameters in southern Australia based on passive acoustics. *Deep Sea Research Part II: Topical Studies in Oceanography*, 157-158:154–168.
- Mering, E. and Chambers, C. (2014). Thinking outside the box: A review of artificial roosts for bats. *Wildlife Society Bulletin*, 38:741–751.
- Meyer, Y. (1990). *Ondelettes et opérateurs: Ondelettes*. Hermann.
- Mikula, P., Morelli, F., Lučan, R., Jones, D., and Tryjanowski, P. (2016). Bats as prey of diurnal birds: A global perspective. *Mammal Review*, 46:160–174.
- Miller, L. A. and Surlykke, A. (2001). How some insects detect and avoid being eaten by bats: Tactics and countertactics of prey and predator. Evolutionarily speaking, insects have responded to selective pressure from bats with new evasive mechanisms, and these very responses in turn put pressure on bats to improve their tactics. *BioScience*, 51(7):570–581.
- Moreno, C. R., Sánchez, L., Macias, S., and Mora, E. C. (2016). Can echolocation calls of Cuban mormopid bats visualized through a heterodyne system? *Caribbean Journal of Science*, 49:27–37.

- Müller, B., Glösmann, M., Peichl, L., Knop, G. C., Hagemann, C., and Ammermüller, J. (2009). Bat eyes have ultraviolet-sensitive cone photoreceptors. *PLOS ONE*, 4(7):1–7.
- Müller, R. (2004). A numerical study of the role of the tragus in the big brown bat. *The Journal of the Acoustical Society of America*, 116(6):3701–3712.
- Neuweiler, G. (2000). *The Biology of Bats*. Oxford University Press.
- Nguyen, B., Morell, C., and De Baets, B. (2017). Supervised distance metric learning through maximization of the Jeffrey divergence. *Pattern Recognition*, 64:215–225.
- Niermann, I., Biedermann, M., Bogdanowicz, W., Brinkmann, R., Le Bris, Y., Ciechanowski, M., Dietz, C., Dietz, I., Estók, P., von Helversen, O., Le Houédec, A., Paksuz, S., Petrov, B., Özkan, B., Piksa, K., Rachwald, A., Y. Roué, S., Sachanowicz, K., Schorcht, W., and Mayer, F. (2007). Biogeography of the recently described *Myotis alcaethoe von Helversen and Heller, 2001*. *Acta Chiropterologica*, 9:361–378.
- Nowak, R. M. (1999). *Walker's mammals of the world*. Baltimore : Johns Hopkins University Press, 6th edition.
- Nyssen, P., Smits, Q., Van De Sijpe, M., Vandendriessche, B., Halfmaerten, D., and Dekeukeleire, D. (2015). First records of *Myotis alcaethoe von Helversen & Heller, 2001* in Belgium. *Belgian Journal of Zoology*, 145:130–136.
- Obrist Martin K., Boesch Ruedi, F. P. F. (2007). Variability in echolocation call design of 26 Swiss bat species: consequences, limits and options for automated field identification with a synergetic pattern recognition approach. *Mammalia*, 68:307–322.
- O'Mara, M. T., Wikelski, M., and Dechmann, D. K. (2014). 50 years of bat tracking: device attachment and future directions. *Methods in Ecology and Evolution*, 5(4):311–319.
- Palmer, J. M., Drees, K. P., Foster, J. T., and Lindner, D. L. (2018). Extreme sensitivity to ultraviolet light in the fungal pathogen causing white-nose syndrome of bats. *Nature Communications*, 9:35.
- Prothero, D. R. (2017). *The Princeton Field Guide to Prehistoric Mammals*. Princeton University Press.
- Root, R. B. (1967). The niche exploitation pattern of the blue-gray gnatcatcher. *Ecological Monographs*, 37(4):317–350.
- Sato, K., Sakamoto, K., Watanuki, Y., Takahashi, A., Katsumata, N., Bost, C., and Weimerskirch, H. (2009). Scaling of soaring seabirds and implications for flight abilities of giant pterosaurs. *PloS one*, 4:e5400.
- Schafer, R. M., Truax, B., and Westerkamp, H. (1973). The Vancouver soundscape.

BIBLIOGRAPHY

- Schnitzler, H.-U. and Kalko, E. K. V. (2001). Echolocation by insect-eating bats: We define four distinct functional groups of bats and find differences in signal structure that correlate with the typical echolocation tasks faced by each group. *BioScience*, 51(7):557–569.
- Schnitzler, H.-U., Moss, C. F., and Denzinger, A. (2003). From spatial orientation to food acquisition in echolocating bats. *Trends in Ecology and Evolution*, 18(8):386–394.
- Senthilkumar, K., Kannan, K., Subramanian, A., and Tanabe, S. (2001). Accumulation of organochlorine pesticides and polychlorinated biphenyls in sediments, aquatic organisms, birds, bird eggs and bat collected from South India. *Environmental Science and Pollution Research*, 8(1):35–47.
- Shannon, C. E. (1949). Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21.
- Simmons, N. B., Seymour, K. L., Habersetzer, J., and Gunnell, G. F. (2008). Primitive early Eocene bat from Wyoming and the evolution of flight and echolocation. *Nature*, 451:818.
- Širović, A. and Hildebrand, J. A. (2011). Using passive acoustics to model blue whale habitat off the western Antarctic peninsula. *Deep Sea Research Part II: Topical Studies in Oceanography*, 58(13):1719–1728.
- Smeeton, N. C. (1985). Early history of the kappa statistic. *Biometrics*, 41(3):795–795.
- Smil, V. (2004). *Enriching the Earth : Fritz Haber, Carl Bosch, and the Transformation of World Food Production / V. Smil*. MIT Press.
- Soady, A. (2013). EU bat rules 'put future of hundreds of churches at risk'. *BBC*.
- Speakman, J. (2001). The evolution of flight and echolocation in bats: Another leap in the dark. *Mammal Review*, 31(2):111–130.
- Speakman, J. and Racey, P. (1991). No cost of echolocation for bats in flight. *Nature*, 350:421–423.
- Speakman, J. R., Anderson, M. E., and Racey, P. A. (1989). The energy cost of echolocation in pipistrelle bats (*Pipistrellus pipistrellus*). *Journal of Comparative Physiology A*, 165(5):679–685.
- Springer, M. S., Teeling, E. C., Madsen, O., Stanhope, M. J., and de Jong, W. W. (2001). Integrated fossil and molecular data reconstruct bat echolocation. *Proceedings of the National Academy of Sciences*, 98(11):6241–6246.
- Stone, E., Harris, S., and Jones, G. (2015). Impacts of artificial lighting on bats: A review of challenges and solutions. *Mammalian Biology*, 80:213–219.
- Sueur, J. and Farina, A. (2015). Ecoacoustics: the ecological investigation and interpretation of environmental sound. *Biosemiotics*, 8(3):493–502.

- Surlykke, A. and Kalko, E. K. V. (2008). Echolocating bats cry out loud to detect their prey. *PLOS ONE*, 3(4):1–10.
- Suzuki, S. and Abe, K. (1985). Topological structural analysis of digitized binary images by border following. *Computer Vision, Graphics, and Image Processing*, 30(1):32–46.
- Teeling, E. C. (2009). Hear, hear: the convergent evolution of echolocation in bats? *Trends in Ecology & Evolution*, 24(7):351–354.
- Tian, L.-X., Pan, Y.-X., Metzner, W., Zhang, J.-S., and Zhang, B.-F. (2015). Bats respond to very weak magnetic fields. *PLoS One*, 10:1–11.
- Ulloa, J. S., Aubin, T., Llusia, D., Bouveyron, C., and Sueur, J. (2018). Estimating animal acoustic diversity in tropical environments using unsupervised multiresolution analysis. *Ecological Indicators*, 90:346–355.
- Vanheuverbeke, N., Boers, K., De Bruyn, L., and Gyselings, R. (2018). Natuurpunt meet temperatuur en luchtvochtigheid in forten om overwinterende vleermuizen beter te beschermen. *Natuurpunt*.
- Voigt, C. C. and Lewanzik, D. (2011). Trapped in the darkness of the night: thermal and energetic constraints of daylight flight in bats. *Proceedings of the Royal Society B: Biological Sciences*, 278(1716):2311–2317.
- Wang, Z. and Bovik, A. C. (2009). Mean squared error: Love it or leave it? A new look at signal fidelity measures. *IEEE Signal Processing Magazine*, 26(1):98–117.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612.
- Warren, R. D. and Witter, M. S. (2002). Monitoring trends in bat populations through roost surveys: methods and data from *Rhinolophus hipposideros*. *Biological Conservation*, 105(2):255–261.
- Webb, P. I., Speakman, J. R., and Racey, P. A. (1996). How hot is a hibernaculum? A review of the temperatures at which bats hibernate. *Canadian Journal of Zoology*, 74(4):761–765.
- Welbergen, J., Booth, C., and Martin, J. (2014). Killer climate: tens of thousands of flying foxes dead in a day. *The conversation*.
- Wilkinson, G. S. and South, J. M. (2002). Life history, ecology and longevity in bats. *Aging Cell*, 1(2):124–131.
- Wilson, D. E. & Reeder, D. M. (2005). *Mammal species of the world : a taxonomic and geographic reference*. Johns Hopkins University Press, 3rd edition.