Probabilistic analysis of coughs in pigs to diagnose respiratory infections

Anna-Lena Giesert*, Wolf-Tilo Balke*, and Gerhard Jahns**

Abstract

In Germany alone, respiratory infections in pigs fattening farms cause losses of millions of Euros. These economic losses are attributed less to total losses due to the death of pigs, but much more to costs for treatment and indirect costs caused by diminished mast und breeding results. The earlier infected animals are detected, the more promising treatments and means to reduce the risk of infection are. Therefore it is necessary to distinguish between coughs from healthy pigs, simple throat cleaning, and coughs by infected pigs. This paper shows that this is automatically possible, and that even the kind of infection can be distinguished. Methods and procedures applied are mainly adapted from those which successfully have been proved in human speech recognition.

Keywords: Pigs, livestock breeding; respiratory diseases; bioacoustics; biosemiotic; signal analysis; Hidden Markov Models (HMM)

Zusammenfassung

Probabilistische Analyse von Hustengeräuschen zur Diagnose von Atemwegserkrankungen bei Schweinen

Infektiöse Atemwegserkrankungen führen allein in Deutschland zu Millionenverlusten in der Schweinezucht. Diese ökonomischen Einbußen sind weniger durch Totalverluste infolge des Verendens der Tiere bedingt, sondern resultieren aus indirekten Verlusten durch die Beeinträchtigungen des Mast- und Zuchterfolges. Je früher erkrankte Tiere erkannt werden, umso erfolgversprechender sind der Behandlungserfolg und Maßnahmen zur Verminderung der Ansteckungsgefahr. Dazu ist es aber notwendig, das Husten eines gesunden Tieres von dem eines erkrankten zu unterscheiden. Der Beitrag zeigt, dass dies automatisch möglich ist, ja dass sich sogar die Art der Infektion erkennen lässt. Dabei kommen, entsprechend angepasst, Methoden und Verfahren zum Einsatz, wie sie sich in der Spracherkennung bewährt haben.

Schlüsselwörter: Schweine, Tierhaltung, Atemwegserkrankungen, Bioakustik, Biosemiotik, Signalanalyse, Hidden-Markov-Modelle (HMM)

---

*   Technische Universität Braunschweig, Institute for Information Systems, Mühlenpfordstraße 23, 38106 Braunschweig, Germany, Email: balke@ifis.cs.tu-bs.de

**   Johann Heinrich von Thünen Institute (vTI), Federal Research Institute for Rural Areas, Forestry and Fisheries, Institute of Agricultural Technology and Biosystems Engineering, Bundesallee 50, 38116 Braunschweig, Germany,

1 Introduction

Respiratory infections in pig fattening farms due to *Proto-bacteria* of the *Pasteurellaceae* family (*Actinobacillus* spec. and *Pasteurella* spec.) cause losses of millions of Dollars or Euros (Scheidt, 1993), beside aspects of animal welfare. These economic losses are attributed less to total losses due to death of pigs, but much more to costs for treatment and indirect costs caused by diminished weight gain and breeding results (Hoy, 1994; Arbeitsgruppe Schwein, 2004). Farmers are reluctant to apply antibiotics as a precaution, not only because of the costs but also because of the risk that the final product, the meat, may contain antibiotic residues.

However, respiratory infections are fast-spreading diseases, which demand treatment as early as possible and fast action to prevent further infections in the herd. This means, the earlier an infection is detected, the better the chances of successful curing and containment of the infection are. In general cough may serve as an early clinical sign of a pig's respiratory tract infection.

The objective is, therefore, to develop software capable of automatically detecting coughing and of distinguishing between morbid coughs and coughs of healthy pigs. The latter also cough from time to time because of inhaled dust, as a kind of throat cleaning. This software would be a basis for a system to monitor a stable as a whole. In general, acoustic monitoring has several advantages in animal breeding: It is neither invasive nor is physical contact necessary. It does not disturb the normal behaviour of the animals and allows continuous monitoring of many animals. The costs are negligible. The software presented here makes use of methods which have proved successful in human speech recognition. The primary focus was to distinguish between morbid and healthy coughs, assuming that the first would effect on the whole respiratory tract, while the latter mainly affects the upper part, thus resulting in generally distinguishable sounds. However, during our experiments it turned out that it is even possible to distinguish between certain kinds of common infections, namely *Pasteurella* and *Actinobacillus*. Obviously, reliable automatically generated evidence pointing to certain infections is a valuable step in disease control, because the respective pigs can be tested and medicated faster. Thus, economic losses in sound monitored facilities could be reduced.

2 Material and methods

2.1 Data

The recorded and labelled data were provided by the courtesy of the Department of Biosystems, Measure, Model & Manage Bioresponses (M3-BIORES), Catholic University of Leuven in Belgium. The recordings contain three different manually labelled types of coughs. Data of the first type are coughs recorded from pigs; derived from a hybrid commercial strain (Landrace × Large White + Danish Duroc boar) at the beginning of the fattening period (100 days old and approx. 40 kg lbw), infected with *Pasteurella multocida*. Records from pigs, suffering from Actinobacillus pleuropneumoniae (three months old and 26 to 35 kg lbw), belonged to a hybrid line by a cross between Italian Landrace × Large White × Duroc. The first and second types of coughs were made under field conditions, in stables. (For more details see Ferrari et al., 2008) The artificially generated coughs were recorded under laboratory conditions, to guaranty, that the pigs (Belgian Landrace of both sexes about 15 kg) were really healthy. Citric acid was evaporated to elicit coughs by irritating the respiratory tract of healthy pigs (throat cleaning). (For more details see Moreaux et al., 1999) Already here, it has to be mentioned, that this inhomogeneity may result in an impairment of the significance of the results.

For training, validating and final testing, 232 *Pasteurella* coughs, 160 *Actinobacillus* coughs, and 149 artificially generated coughs were available. 70 % of each data type was used for training, whereas 30 % was used for validation of the recognition rate.

2.2 Data cleaning

As stated above, the morbid cough recordings contain ambient noises, since they were recorded in a stable. In the case of *Pasteurella* and *Actinobacillus* infections, these background noises were similar. But the artificially induced coughs lack these typical stable noises. This holds the risk of misclassification, because the classifier may be trained more on background noises than on cough types. Therefore the ambient noise had to be reduced as much as possible. Especially the sounds of ventilators or rattling chains are characterized by high frequencies, in contrast to the dominant low frequencies of coughs. Therefore we used a low pass filter to cut off high frequencies. Comparing filtered copies and unfiltered original data sets for training and evaluation purposes it turned out that the best cut-off frequency is 20 kHz, about 0.45 of the normalized Nyquist frequency of each signal. However, it must be mentioned that this analysis depends on human hearing. This implies that it cannot be guaranteed, that this cut-off frequency is optimal under all circumstances for the statistical models used later. In any case, as a result of this low pass filter, manual inspection showed clear coughs with only minimal ambient noise in most signals of the test set.

2.3 Feature extraction

Before coughs can be automatically classified, the significant characteristics of individual pig coughs have to be extracted from the recordings. We used a set of five well-known feature extraction methods and also combined their feature sets. Since all these methods only work on discrete spectra of the signals, we split each continuous signal into a set of equidistant frames. Each frame has a length of 11 ms, such that each frame can be regarded as a discrete stationary observation. Errors may occur when splitting the signal into frames, especially on the borders between frames (Scheithauer, 2005). Therefore we applied a Hamming Window to each frame to reduce such errors. Furthermore, all frames are slightly overlapped to guarantee smooth transitions. After discretization, the time-based signals had to be transformed to the frequency domain, because all extraction methods used in this work are based on spectra. The most popular algorithm for this is Discrete Fourier Transform (DFT) representing the signal by the integral over a sum of sine and cosine waves (Chirlian, 1994). A second method is the Discrete Cosine Transform (DCT), which works similarly to the DFT, but represents the signal in terms of a sum of cosine waves (Seibt, 2006; Stearns, 2002). We also used DCT as a baseline for our comparisons between feature accuracies.

Coughing sounds are very similar to each other, but are always strongly influenced by the pigs' vocal tracts. Thus, we used four features specifically designed for distinguishing vocalizations in speech recognition: the first pair is based on cepstral coefficients, the second pair is based on linear predictive coding. All algorithms for extraction were taken from the seminal work in Rabiner and Juang (1993). For the first pair we first extracted the spectral energy of a signal, the so-called cepstral coefficients, by applying a logarithmic dampening on the DFT. Then the Mel frequency cepstral coefficients (MFCCs), where all frequencies are mapped to the Mel-Scale in discrete steps such that the frequency bands are reduced, were extracted. A similar extraction was done for the Bark frequency spectral coefficients (BFCCs) with the Bark-Scale instead of the Mel-Scale. The use of both scales is justified due to the similarity of the vocal tract and the auditory system in mammals respecting the Weber-Fechner Law. The third feature used are the coefficients of linear predictive coding (LPC), which extracts typical speech parameters like pitch and vocal tract area. In particular LPC shows the degree of self-similarity of a signal, i.e., any frame can be represented as a linear combination of a fixed number of preceding frames, where all coefficients are assumed to be constant. As a last feature, the coefficients of perceptual linear predictive coding (PLPC) were extracted. In contrast to LPC, which is strongly dependent on the individual speaker, PLPC is speaker, respectively animal, independent. Basically, here all frequencies are first warped to the Bark-scale before extracting LPC coefficients (Hermansky, 1990).

Finally the first 20 coefficients of DCT, BFCC, MFCC, LPC and PLPC, respectively, were used as a feature vector. Each cough frame is represented by a vector containing all five features (i.e., a 100-dimensional feature vector) and each cough recording is represented as a concatenation of several feature vectors depending on its respective length.

2.4 Classification

The actual classification part was derived using Hidden Markov Models (HMMs), which have proven to be useful in speech recognition and also in animal call recognition, like calls from cows (Jahns, 2007), or birds (Wolff, 2008). Moreover, HMMs allow the time-independent modelling of signals and the differentiation of variations in the spectra for recognition independent of the individual speaker/ vocalizing animal.

Generally speaking, a Hidden Markov Model is a doubly embedded stochastic process, where the states of the model are hidden (for more detailed information see e.g. Rabiner, 1989). Hidden Markov Models are typically used where the lengths of feature vectors vary from one to another. A Hidden Markov Model is characterized by: a set of distinguishable states, a matrix of transition probabilities from one state to another, a vector of observation symbol probabilities for each state, which is needed for the probability that a state generates an output and an initial state distribution. The challenge is to train a generic HMM to obtain the best fitting of these statistical parameters by samples. In our case the number of samples is the number of labelled coughs available. Because it is a statistical process, it is evident that the number of samples should be as great as possible. Each HMM has to be trained: in particular, the model parameters of the HMM are estimated by applying the well-known Baum-Welch Algorithm. The likelihood that the HMM has produced some observation (in our case a certain set of feature vectors representing a recording) can be determined using the Viterbi Algorithm, basically a linear programming algorithm to determine the most probable sequence of states for a given observation.

We used a simple left-right HMM for modelling cough sounds and inspected several spectrograms to get an intuition about the necessary number of states. Unfortunately there are no general rules to determine the number of states. So we used common heuristics: a first assumption was made by dividing cough sounds into at least three sections (heavy inhaling, strong cough, ebbing away) and allowing a transition from one state to the next, remaining in one state or skip one state, as shown in Figure 1. In our experiments we then varied the number of states

from three to seven to check for the most suitable number of states showing the highest recognition rates. For each type of cough (*Pasteurella*, *Actinobacillus*, and non-morbid cough) an individual HMM was trained. To classify unknown cough sounds, the most likely HMM (i.e., the one best explaining the complete chain of observations) has to be computed. This procedure results in a three-way classifier whose correctness can be directly measured in the number of false classifications after evaluating all three hypotheses.
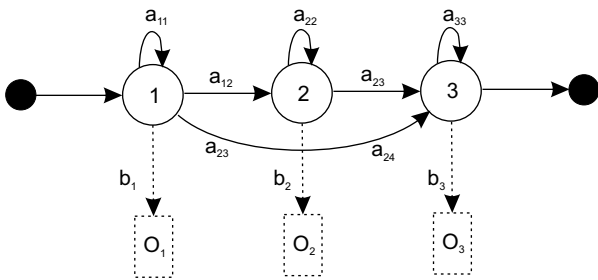
Figure 1:

A left-right HMM with three states. $\lambda = (\pi, A, B)$ initial state distribution (π), state transition probabilities (A), observation probability distribution (B).

3 Results

3.1 Total recognition rate

The experiments showed recognition rates for unknown recordings of more than 97 %. In particular, the HMMs trained with the original unfiltered data set reached 100 % recognition rate for *Pasteurella* coughs, 97.3 % for the *Actinobacillus* coughs and 98.4 % for the artificially generated coughs. Moreover, the representation of the states of the HMM was very good. The HMMs trained with the filtered data set reached even better recognition results than the models for the unfiltered data set, however, the representation of the states of this HMM was worse. It reaches 99 % recognition rate for the *Pasteurella* coughs, 100 % for the *Actinobacillus* coughs and 98.4 % for the artificially elicited coughs.

In summary both experiments, filtered and unfiltered, have shown that it is possible to accurately distinguish between artificially induced and morbid coughs, as well as between the two morbid cough types from pigs with *Pasteurella* or *Actinobacillus* infections.

3.2 Feature extraction

The experiments also confirmed that extraction methods designed for speech recognition are beneficial for cough classification. All individual features performed sig-

nificantly better than general frequency-based features. DCT- features returned entirely unacceptable recognition rates of around 35 % (i.e., slightly better than random guessing) as shown in Figure 2. Still, the average recognition rate combining all features (including DCT) in one classifier was slightly higher than all other combinations of features. Thus, in combination with the other features the DCT feature still contributes to the excellent classification performance shown in Figure 3.
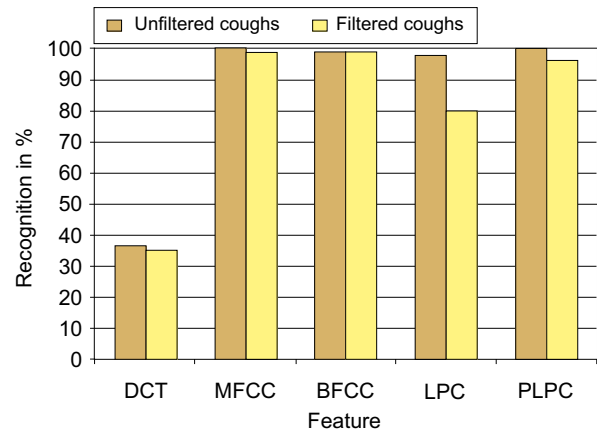

Figure 2:

Average percentage of correctly classified coughs with five different feature extraction methods
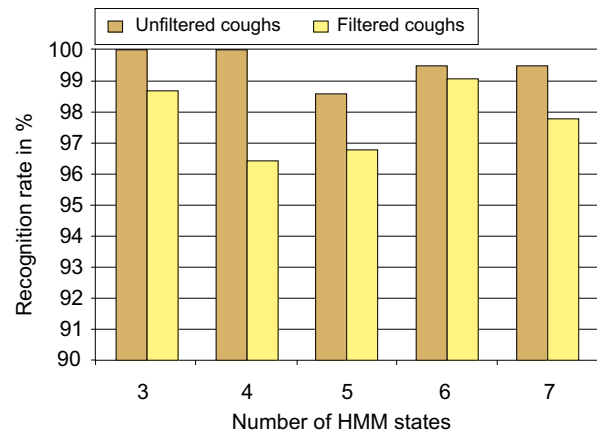

Figure 3:

Average recognition rate of filtered and unfiltered coughs for each tested HMM

3.3 Suitability of the Hidden Markov Models

The experiments showed that HMM with five states perform best for unfiltered data (see example classifications in Figure 4, 5 and HMM with six states for filtered data in Figure 6). Therefore, not only the recognition rate, but also the state representation is a crucial factor for finding the best HMM. For instance, a HMM for the unfiltered data

set with a total recognition rate of 100 % and three states can be defined (see Figure 1). However, this first impression of perfect recognition may be misleading, because the state representation is entirely insufficient. As can be seen in Figure 4, coughs are always meaningfully divided into the different states. Beyond optimizing the recognition rate, it is important, especially for the model validity and extensibility, that this state representation is the same for all coughs, which indeed has similar reasons. Hence, some HMM (even with slightly better recognition rates over the test set) are still not suitable for a good classification because of their state representation. The recognition rates for all tested HMM can be seen in Figure 3. The few erroneously recognised unfiltered coughs were two *Pasteurella* coughs which were classified as *Actionbacillus* coughs and one artificially generated cough which was classified as *Actinobacillus* cough. Further, one filtered *Pasteurella* cough was classified as *Actionbacillus* cough and one filtered artificially generated cough was classified as *Actionbacillus* cough.
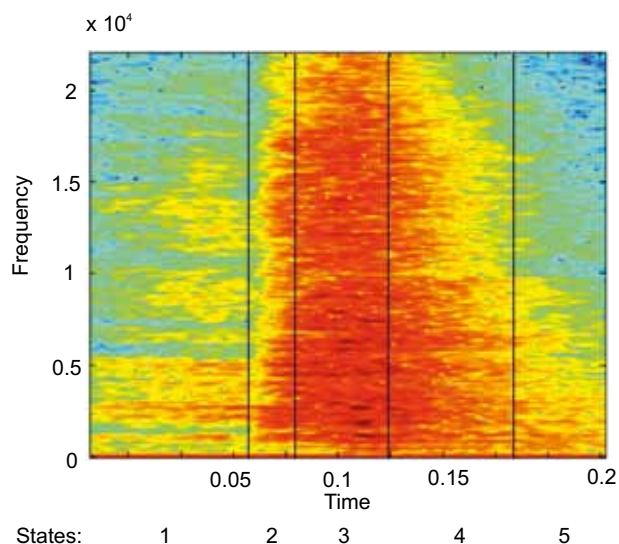

Figure 4:
Representation of five HMM states for a cough of a healthy pig elicited a by citric acid

Moreover, Figure 5 shows that not every cough has to be represented by the entire sequence of states. This depends on the differences of coughs within each cough type and adapts perfectly to these differences. For example, *Actinobacillus* coughs exist which have much higher frequencies at the beginning of the cough and more intensity than the cough shown in the Figure. This feature is modelled by a second alternative state, which is skipped for this cough.
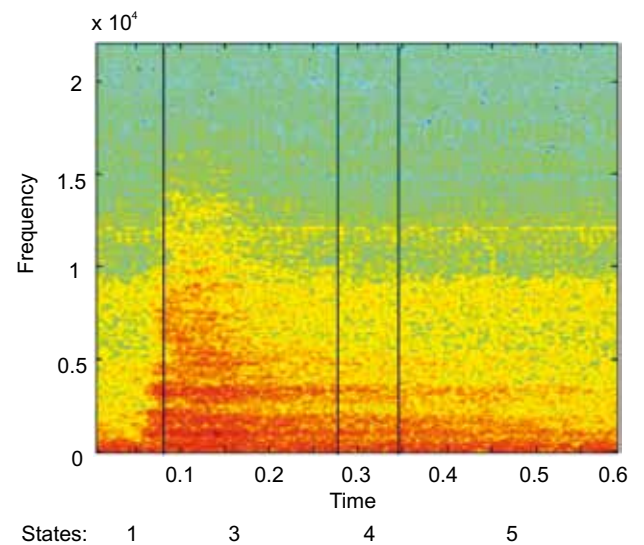

Figure 5:
Representation of the five states HMM for an unfiltered recording of *Actinobacillus* cough. Note: state two is skipped in this example. This is directly the desired effect of the modelled simple left-right HMM.
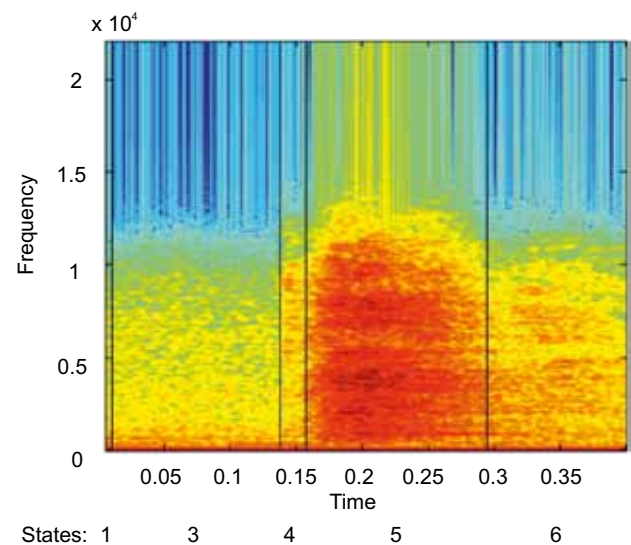

Figure 6:
Representation of a six states HMM for a filtered cough of a healthy pig elicited a by citric acid

Finally, short computation times are necessary for the usage of our software as an online recognition system. Of course the classification time always depends on the length of the cough type, the feature extraction methods and the classification method. All our experiments were run on a 2.4 GHz personal computer. Thus the computation time needed to classify a single cough (including feature extraction) is short enough, such that the classification can be performed in real time.

4 Discussion

Our results show that the classification with HMMs in combination with feature extraction methods from the area of speech recognition works fine. However, there are still several problems to be discussed:

- From speech recognition it is known, that LPC is sensitive to characteristics of the individual vocal tract (age, gender etc.) while especially MFCC and BFCC are not. This explains the better performance of the latter (s. Figure 2).
- The total recognition rate and state representation of the HMM show that data cleaning does not imply better results. This depends, on one hand, on our chosen filter method, and on the other on variances of the frequency bands. Having only a single fixed cut-off frequency for all cough types was shown to reduce the recognition rate, i.e., a simple low pass filter will sometimes also obviously destroy information important for correct classifications.
- While the chosen feature extraction methods work well, there are still two disadvantages. One problem is that these feature extraction methods are all based on frequency spectra which could be a disadvantage for fault-tolerant classification when extending the classification system. Furthermore the proposed feature extraction is complex to compute. Simpler time-based extraction methods might reach comparable classification accuracy with the advantage of lower computation time.
- A significant problem is also the small number of data available. More data (either recorded in an identical environment or coughs of the same type recorded in different environments) is needed to ensure accuracy, robustness and reliability.
- And finally, all experiments conducted here were performed on manually cut and labelled recordings of coughing events. However, to automatically monitor breeding facilities cough sounds have to be spotted within of a continuous sound stream, despite all of the ambient noise in a stable. To spot coughs in continuous recordings buried in noise will be the next problem to be addressed.

5 Conclusions

This paper explains how to classify coughing sounds of pigs using Hidden Markov Models. Focusing on *Pasteurella* and *Actinobacillus* infections, as compared to artificially induced coughs (e.g. induced by dust), we investigated the necessary data cleaning, feature extraction and model training. The results show that the classification with HMMs in combination with feature extraction methods from the area of speech recognition works excellently. Also, classification times even for such complex computations are satisfactory. Thus, this method is definitely acceptable for classifying pig coughs using a monitoring software.

References

Arbeitsgruppe Schwein (2004) Programm zur Überwachung und Bekämpfung der progressiven Rhinitis atrophicans bei Zuchtschweinen. Amtl Veterinärnachr 77(8)8:97-113

Chirlian PM (1994) Signals and filters. New York : Van Nostrand Reinhold, 397 p

Ferrari S, Silva M, Guarino M, Berckmans D (2008) Analysis of cough sound for diagnosis of respiration infections in intensive pig farming. Trans Am Soc Agric Biol Eng 51(3):1051-1055

Hermansky H (1990) Preceptual linear predictive (plp) analysis of speech. J Acoust Soc Am 87(4):1738-1752

Hoy S (1994) Zu den Auswirkungen von Atemwegserkrankungen auf die Mast- und Fruchtbarkeitsleistungen der Schweine. Prakt Tierarzt (2):121-127

Jahns G (2008) Call recognition to identify cow conditions : a call-recogniser translating calls to text. Comput Electron Agric 62(1):54-58

Moreaux B, Beerens D, Gustin P (1999) Development of a cough induction test in pigs : effects of SR 48968 and enalapril. J Vet Pharmacol Ther (22):387-389

Murphy K (2005) Hidden Markov Model (HMM) toolbox for Matlab [online]. To be found at <http://people.cs.ubc.ca/~murphyk/Software/HMM/hmm.htm> [quoted 31.05.2011]

Rabiner L (1989) A tutorial on hidden Markov models and selected applications in speech recognition. Proceedings IEEE 77(2):257-286

Rabiner L, Juang B-H (1993) Fundamentals of speech recognition. Upper Saddle River NJ : Prentice Hall PTR, 507 p

Seibt P (2006) Algorithmic information theory : mathematics of digital information processing. Berlin : Springer, 442 p

Scheidt A (1993) Mycoplasmal pneumonia [online]. To be found at <http://www.ncsu.edu/project/swine_extension/healthyhogs/book1993/scheidt1.htm> [quoted 31.05.2011]

Scheithauer R (2005) Signale und Systeme : Grundlagen für die Mess- und Regelungstechnik und Nachrichtentechnik. Stuttgart : Teubner, 436 p

Stearns SD (2003) Digital signal processing with examples in MATLAB. Boca Raton : CRC Pr, 336 p

Wolff D (2008) Detecting bird sounds via periodic structures : a robust pattern recognition approach to unsupervised animal monitoring [online]. To be found at <http://www-mmdb.iai.uni-bonn.de/download/Diplomarbeiten/Diplomarbeit_Daniel_Wolff.pdf> [quoted 01.06.2011]