Developing a Bayesian Network for Clinical Diagnosis in Veterinary Medicine: from the Individual to the Herd

Petra L. Geenen and **Linda C. van der Gaag**
Department of Information and Computing Sciences, Utrecht University,
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
e-mail: {petrag,linda}@cs.uu.nl

Abstract

In close collaboration with two experts, we are building a Bayesian network for the early detection of classical swine fever in pigs. Although our network will ultimately be tailored to detecting the disease in pig herds, we decided to begin with modelling the pathogenesis of the disease in the individual animal. Having modelled the patterns of classical swine fever observed in individuals, we now are in the process of rendering our network applicable to herds. Upon doing so, we are confronted with various intriguing modelling issues. As an example, we focus in this paper on the issue of modelling diverging knowledge that pertains to the different types of pig in a herd.

1 Introduction

Bayesian networks are being successfully applied in a wide range of domains, for a variety of problems. Most notably, networks are being developed for the medical domain, for solving diagnostic problems. In the medical domain, diagnosis is aimed at detecting disease in an individual patient. Bayesian networks for clinical diagnosis therefore typically model the onset and progression of a disease in an individual. Over the last decade in fact, considerable experience has been gained in modelling the pathogenesis of disease [1].

Although the veterinary domain at first sight is closely related to the medical domain of human disease, veterinary applications of Bayesian networks are still quite rare, although recently some networks have been proposed, for example for diagnosing bovine diseases [2] and for establishing the severity of a specific infection in swine [3]. A major difference in diagnosis in the veterinary field compared to that in the human medical domain, is that the detection of disease is focused on herds rather than on individual animals, especially when infectious diseases in farm animals are concerned. Developing a Bayesian network for the veterinary domain therefore demands that upon modelling the pathogenesis of an infectious disease, its onset and progression within a herd needs to be taken into account. While considerable experience has been gained with building Bayesian networks for the detection of disease in individuals, modelling the knowledge involved in detecting diseases in a population has so far received little attention.

In close collaboration with two experts from the Central Institute of Animal Disease Control in the Netherlands, we are developing a Bayesian network for the early detection of classical swine fever in pig herds. Classical swine fever is an infectious viral disease of pigs, which has serious socio-economical consequences upon an outbreak. As the disease has a potential for rapid spread, it is imperative that its occurrence is detected in the early stages. The Bayesian network under construction is aimed at supporting veterinary practitioners in the diagnosis of the disease when visiting pig farms with disease problems of unknown cause.

Although our network for classical swine fever will ultimately be tailored to detecting the disease in pig herds, we decided to begin with modelling the pathogenesis of the disease in individual animals. By doing so, we could carefully model the details of the onset and progression of the disease within an animal, without being concerned with issues related to the transmission of the disease. In addition, we could build upon our experience with modelling the pathogenesis of human disease in individual patients. Having modelled the patterns of classical swine fever observed in individual animals, we now are in the process of rendering our network applicable to herds. Upon doing so, we are confronted with various intriguing modelling issues.

In this paper, we address one of the modelling issues that we have encountered so far upon tailoring our network for classical swine fever to pig herds. We will focus more specifically on the issue of modelling diverging knowledge that pertains to the different types of pig in a herd. Upon studying the pathogenesis of the disease, we noted that the progression of the disease and the clinical signs to be observed, differ among the various types of pig in a herd. More in particular, in inseminated sows the disease may

lead to an intra-uterine infection, which may give rise to an early abortion or to a litter with a relatively large number of stillborn or trembling piglets. It will be evident that the part of the network that pertains to the reproductive cycle should be taken to apply to sows only. To model this diverging knowledge, we enhanced the meanings of the variables that capture issues related to the reproductive cycle in sows, by including a meta-level value that serves to represent applicability of the variable. To provide for correct probabilistic reasoning, we formulated associated parameter probabilities such that the variables involved partake in the reasoning processes only for sows.

The paper is organised as follows. In Section 2, we briefly introduce the domain of classical swine fever and the network that we developed so far. In Section 3, we elaborate on both the added value and the restrictions that originate from combining knowledge of the disease in individual animals and at the population level for the purpose of early detection. In Section 4, we focus on the modelling of diverging knowledge for different types of pig. The paper ends with our concluding observations in Section 5.

2 The network for classical swine fever

Classical swine fever is a viral disease of pigs with a potential for rapid spread. The virus causing the disease is transmitted mainly by direct contact between pigs, yet transmission by for example farmers or vehicles is also known to occur. When a pig is infected, the virus first invades the lymphatic system. It subsequently affects the blood vessels and the immune system, which may give rise to bleedings and diminished resistance to secondary infections. The virus will ultimately affect several organs and the pig will die. As a consequence of the infection, a pig will show different disease symptoms, among which are fever, inflammation of the eyes, neurological disorders, and haemorrhages of the skin.

Classical swine fever is quite common in parts of Europe and Africa, and in many countries of Asia and of Central and South America. Extensive measures have been taken within the European pig husbandry to prevent the introduction and spread of the virus. Unfortunately, however, each year several outbreaks occur. Since an outbreak of the disease has a major impact on international trade of animals and animal products, veterinarians are obliged by law to report any suspicion of classical swine fever in a pig herd.

Clinical symptoms seen by the farmer or by the veterinarian are usually the first indications of the presence of classical swine fever in a herd. Unfortunately, the symptoms of the disease are rather atypical and are shared to a large extent by common airways and gastro-intestinal infections. As a consequence, the disease can remain undetected for weeks and may spread to many herds. In the 1997/1998 epidemic in the Netherlands, for example, it was estimated that the disease remained undetected for six weeks and that, by that time, already 39 herds were infected [4]. This major outbreak had serious economical consequences. A total of 12 million pigs had to be killed and the costs involved were estimated to be 2.3 billion US dollars.

Because of the major socio-economical consequences that an outbreak may have, reducing the time between first infection of a herd and first detection is of major importance. The shorter this high-risk period, the more restricted the epidemic may be. We feel that improving the clinical diagnosis of the disease, that is, its detection based upon clinical signs seen at the farm, is an important step towards shortening the high-risk period. In close collaboration with two experts from the Central Institute of Animal Disease Control in the Netherlands, we are building a Bayesian network that is tailored to early detection of the disease; the network is aimed at usage by veterinary practitioners in the field.

Our network is still under development and currently includes 42 variables for which over 2400 parameter probabilities have been assessed. The variables in the network model the risk factors and the pathogenesis of the disease, but more specifically the network also models the clinical signs to be observed in a pig herd to provide for diagnosis at a farm site. In the one and a half years since the beginning of the construction of the network, we held one unstructured interview in which the experts were asked to describe the domain and 11 structured interviews in which the experts were asked detailed questions. In six of these structured interviews, the probabilities required for the network were obtained using standardised forms with questions accompanied by a probability scale containing words and numbers [5]. In all sessions, both experts were present and consensus was always reached.

3 From the individual to the herd

To provide for the diagnosis of classical swine fever, a Bayesian network should include not just knowledge about the pathogenesis of the disease in individual animals, but also knowledge about its progression in a herd. We briefly describe the sequence of events during an outbreak and comment on the modelling issues that arise upon capturing the knowledge involved.

The first pigs that are infected with classical swine fever on a farm will not immediately show clinical signs of the disease, as a consequence of an incubation period of 2 to 14 days. After the incubation period, these pigs will start to develop clinical signs, but the severity of these signs and, hence, their visibility will depend on many factors like the strain of the virus and the pig's immune status. At this stage of the disease, the pigs are infectious and, since the virus rapidly spreads in a susceptible pig population, the fraction of infected pigs in the herd will increase over time. Depending on the attentiveness and experience of the farmer

and possibly the veterinarian involved, signs of the disease will sooner or later be noticed and a suspicion of classical swine fever may be reported.

In the early diagnosis of a herd with disease problems of an unknown cause, a veterinarian will at first focus her attention on individual pigs. A combination of clinical signs in a single pig, such as fever and inflammation of the eyes, then is more likely to give rise to a suspicion of classical swine fever than if these signs are observed separately in different pigs. Diagnostic reasoning therefore explicitly begins by focusing on individual pigs. Upon finding a diseased pig, however, the veterinarian will consider the other pigs in the same pen. In case more pigs show signs of disease, an infectious cause becomes a more and more likely explanation of the observed findings. The additional information from the pen thus is drawn into the diagnostic process. Similarly, information from across pens, such as an increased number of sows with stillborn piglets, can provide further evidence of an infectious disease.

The progression of classical swine fever in an individual pig is dynamic, that is, during the infection different clinical signs are found to succeed one another. In an infected pig, for example, constipation is usually followed by diarrhoea. Since in an infected herd the various pigs will typically be in different stages of the disease, a mixture of signs will be observed. Moreover, due to the immunodepression caused by the disease, opportunistic pathogens are likely to cause secondary infections. The result is a highly confusing clinical picture of signs. To contribute to the complexity of the diagnosis, the veterinarian can only observe the pigs during her visit to the herd and, although the farmer can inform her about the disease progression within the herd, the provided information generally is not on individual pigs.

As outline above, diagnostic reasoning for the early detection of classical swine fever in a pig herd typically sets out by considering individual animals and then proceeds to an investigation of the pen and of the entire herd. A Bayesian network that is aimed at supporting a veterinarian in her reasoning tasks, therefore, should capture knowledge about the onset and progression of the disease at three different levels of scope. Moreover, since the clinical picture presented to the veterinarian is likely to be a mixture of various simultaneous infections, the network cannot be restricted to just the effects of classical swine fever: it should be able to distinguish between the more common airways and gastrointestinal infections and the low-prevalence classical swine fever. Ideally also the network should be able to handle changing patterns of disease over time. To build a network that meets these requirements clearly is a challenge to the knowledge engineers involved.

Although our network will ultimately be tailored to the early diagnosis of classical swine fever in herds, we decided to begin with modelling the onset and progression of the disease in individual animals. From the considerations above, we have that diagnostic reasoning involves at least reasoning about individual pigs, which requires a model of the disease's pathogenesis. By beginning the construction of our network at the level of individual animals, we could carefully model the intricate details of the progression of the disease, building upon our experience with modelling the pathogenesis of human disease in individual patients. We further decided to develop a static network. Because veterinary routines do not monitor individual pigs over time, the ability to reason about processes that change with time, as provided by a dynamic network, would currently be of little use in practice.

4 Modelling diverging knowledge

Upon constructing our Bayesian network for classical swine fever, we decided, as argued above, to begin with modelling the onset and progression of the disease in individual pigs and use the resulting model for developing a network for the early detection of the disease in pig herds. Upon studying the pathogenesis of the disease, we noted that the progression of the disease and the clinical signs to be observed differ among the various types of pig in a herd. More specifically, in inseminated sows the disease may lead to an intra-uterine infection, which may give rise to an early abortion or to a litter with a relatively large number of stillborn or trembling piglets.

4.1 Tailoring the graphical structure

Since a herd includes various different types of pig and our ultimate aim is to construct a model for detecting classical swine fever in herds, we would like to construct our network for the pathogenesis of the disease to apply to any type of pig. In essence, the framework of *relational Bayesian networks* can be used to capture the diverging knowledge that pertains to the various types of individual in a population among which the specifics of disease may differ; the basic idea then is to capture the knowledge of the various subpopulations in different networks [6]. When modelling the progression of classical swine fever in sows on the one hand and in boars, piglets and gilts on the other hand, the resulting networks would be quite similar, however. The graphical structures of the networks would in fact differ only in the part that pertains to the reproductive cycle in sows, which includes just five variables. The only other differences between the networks would be in the parameter probabilities involved. For example, at a pig farm sows and gilts are housed differently and as a consequence do not have equal probabilities of becoming infected. Once infected, however, the progression of the disease is quite similar for all types of pig. Since the differences between the networks for the various types of pig are so few, we decided not to exploit the framework of relational Bayesian

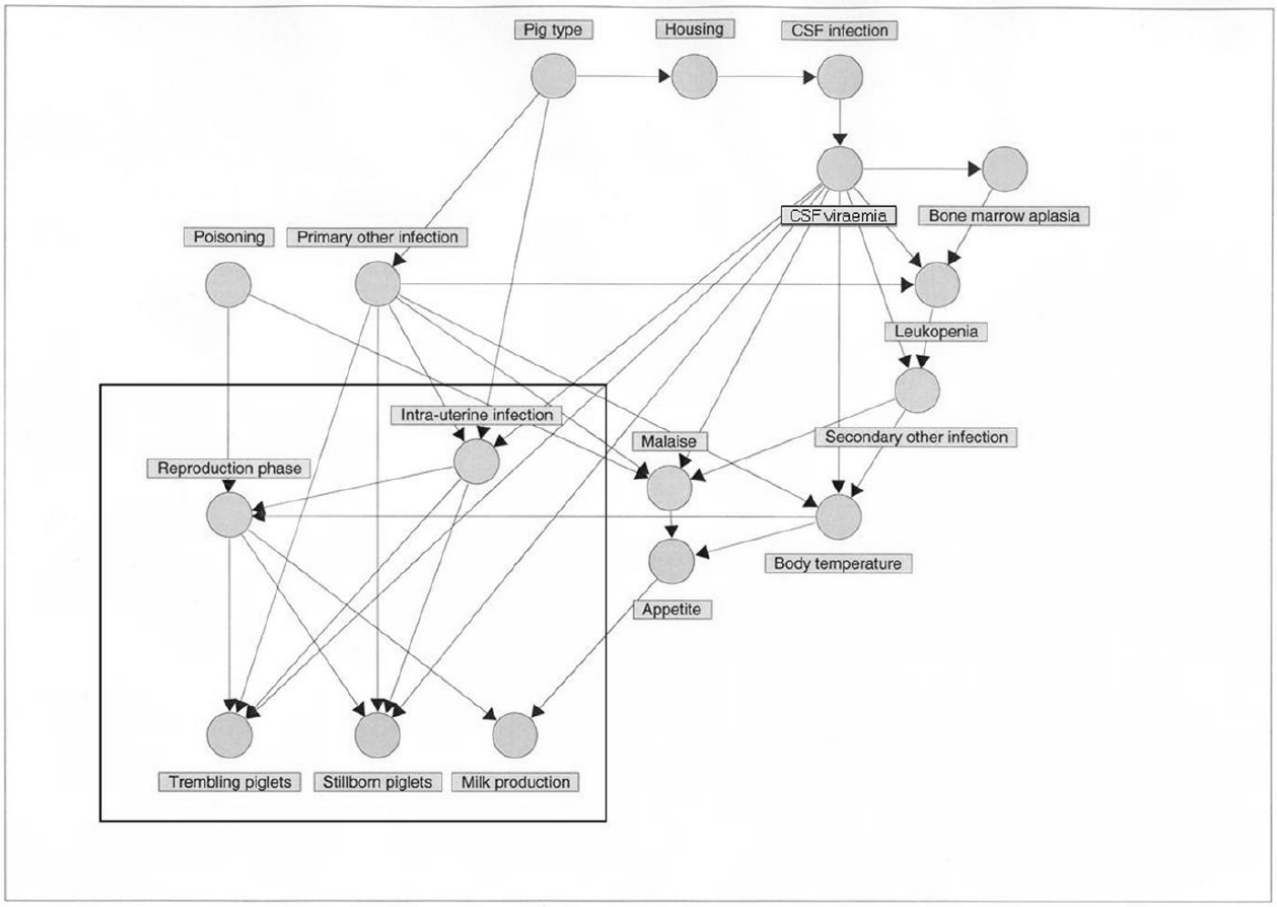

Figure 1: Part of the Bayesian network for classical swine fever.

networks because of the overhead it would generate. We decided to model the pathogenesis of the disease for the different types of pig in a single network instead.

Upon modelling the knowledge about the onset and progression of classical swine fever in individual pigs, we had to pay special attention to the differences in the pathogenesis of the disease for the various types of pig. Some of the issues involved in modelling diverging knowledge in a single network, are quite readily resolved. Most differences in the probabilities of becoming infected with the disease, for example, could be captured by including two additional variables; these are the variable *Pig type*, modelling the various types of pig, and the variable *Housing*, which captures whether pigs are housed individually, in a small group, or in a large group. Correctly modelling the differences between sows and the other types of pig with regard to issues of reproductivity, however, turned out to be more involved.

We consider within our network the part that pertains to the reproductive cycle in sows. For ease of reference, Figure 1 shows part of our overall network; the part that pertains to the reproductive cycle in sows is shown within the indicated box. As an example, we focus on the variable *Intra-uterine infection*. As mentioned before, in sows an infection by classical swine fever may give rise to an intra-uterine infection. We distinguish between infections that occur in the early stages of the pregnancy and infections that occur later in the gestation period, since these tend to give different problems; while an early infection results in an abortion, an infection in the later stages of the pregnancy will result in a relatively large number of stillborn or trembling piglets in the litter. The values of the variable *Intra-uterine infection* therefore are *no*, *early* and *late*.

It will be evident that the variable *Intra-uterine infection* should be taken to apply to sows only. The knowledge that the other types of pig will not develop an intra-uterine infection, can in essence be expressed by the value *no* of the variable. For boars, for example, the value then captures the knowledge that a boar never has an intra-uterine infection. One of the representational problems associated with taking the value *no* for capturing this knowledge, is that the value no longer has a well-defined semantics. Where the value *no* would indicate for a sow that she does not have an intra-uterine infection while in essence she could have one, the same value would indicate for a boar that he does not have such an infection because he can never have one.

The double meaning of the value *no* poses yet another, even more important representational problem. While the value *no* for a boar would not have any weight with respect to the presence or absence of classical swine fever, the value *no* for a sow would provide relevant diagnostic evidence: if a sow does not show any signs of an intra-uterine infection, then this is construed as a contra-indication against an infection by the virus of classical swine fever. By assigning a double meaning to the value *no* as indicated above, no distinction can be made with respect to the value's weight in diagnostic reasoning for different types of pig.

Because of the representational problems outlined above, we decided not to use the value *no* for capturing the knowledge that except for sows the other types of pig will not develop an intra-uterine infection. Instead we decided to add the new value *not applicable* to the collection of values for the variable *Intra-uterine infection*. The additional value in essence captures the knowledge that for other types of pig than sows, the variable has no meaning. Note that the value *not applicable* now differs from the values *no*, *early* and *late* in that it captures meta-knowledge about the variable. Similarly, we included the value *not applicable* in the collections of values for the variables *Reproductive phase*, *Trembling piglets*, *Stillborn piglets* and *Milk production*. Note that the five thus extended variables constitute the part of the network that pertains to the reproductive cycle in sows.

4.2 Detailing the parameter probabilities

To provide for modelling the diverging knowledge about the progression of classical swine fever in different types of pig, we have so far redefined the meanings of the five variables that represent issues related to reproductivity. To correctly capture their new meanings, we now focus on the parameter probabilities for these five variables. We note that these parameter probabilities should serve to effectively preclude the five variables from diagnostic reasoning whenever the network is being consulted for another type of pig than a sow. For example, if the network is used for establishing the presence or absence of classical swine fever in a boar, each of the five variables should adopt the value *not applicable* with certainty. These values, however, should not induce, just by themselves, a dependence between for example the variables *Primary other infection* and *CSF viraemia*. If the network is consulted for a sow, on the other hand, the five variables should partake in the reasoning process. Any evidence about problems with the sow's litter, for example, should then induce dependences between the various possible explanations.

As an example, we consider the parameter probabilities for the variable *Intra-uterine infection*. We recall that, when taken to pertain to a sow, this variable has the three values *no*, *early* and *late*. The probability table for the variable then includes 24 parameter probabilities; these parameter probabilities have been assessed by our two experts and are shown in Table 1 for ease of reference. When taken with respect to all types of pig, the variable has an additional incoming arc from the variable *Pig type*, which captures the five different types of pig distinguished in our domain of application; note that the additional arc provides for explicitly distinguishing between these five types in terms of the probabilities involved. The variable moreover has the additional value *not applicable* in its collection of possible values. When taken for all types of pig, the probability table for the variable includes 160 rather than 24 parameter probabilities.

Table 1: The probability table for the variable *Intra-uterine infection* given the variables *Primary other infection* and *CSF viraemia*.

Intra-uterine infection	*Primary other infection*	*CSF viraemia*	
no	*none*	*yes*	0.25
no	*none*	*no*	1.00
no	*airways*	*yes*	0.2
no	*airways*	*no*	0.99
no	*gastro-intestinal*	*yes*	0.2
no	*gastro-intestinal*	*no*	0.99
no	*airways and gastro*	*yes*	0.14
no	*airways and gastro*	*no*	0.975
early	*none*	*yes*	0.634
early	*none*	*no*	0
early	*airways*	*yes*	0.676
early	*airways*	*no*	0.0084
early	*gastro-intestinal*	*yes*	0.676
early	*gastro-intestinal*	*no*	0.0084
early	*airways and gastro*	*yes*	0.73
early	*airways and gastro*	*no*	0.021
late	*none*	*yes*	0.116
late	*none*	*no*	0
late	*airways*	*yes*	0.124
late	*airways*	*no*	0.0016
late	*gastro-intestinal*	*yes*	0.124
late	*gastro-intestinal*	*no*	0.0016
late	*airways and gastro*	*yes*	0.13
late	*airways and gastro*	*no*	0.004

We observe that, by taking the variable *Intra-uterine infection* to apply to all types of pig, many more parameter probabilities are required for its probability table than for the variable in its meaning for just sows. The domain experts, however, need not provide assessments for these additionally required probabilities. Since the knowledge that has been added does not involve any uncertainty, all additional parameter probabilities equal either 0 or 1. The assessments provided by the experts now are taken for the

parameter probabilities that are conditioned on the pig being a sow. The parameter probabilities for the absence or the presence of an early or a late intra-uterine infection in a boar, for example, are taken to be equal to 0 given any combination of values for the two possible causes of such an infection. The probability of the value *not applicable* given any such combination then equals 1.00. Note that the parameter probabilities with respect to a boar now match the meaning that we intended for the variable and its values.

We recall that also the meanings of the variables *Reproduction phase*, *Stillborn piglets*, *Trembling piglets* and *Milk production* have been extended to pertain to any type of pig. For the probability tables for these variables, similar observations hold as outlined above for the variable *Intra-uterine infection*. Again, the probability tables are readily extended without the need for the domain experts to supply new probability assessments. We would like to note that, unlike the variable *Intra-uterine infection*, the other four variables pertaining to the reproductive cycle in sows do not require an additional incoming arc from the variable *Pig type*. We observe that these variables all include the variable *Intra-uterine infection* among their causes. If, upon reasoning with the network, this latter variable adopts the value *not applicable* with certainty, then this value is transferred directly to the other four variables. For these variables, it is no longer of interest whether the value originated from the pig being a boar, a piglet or a gilt.

To conclude, we address the reasoning behaviour of the part of the network that pertains to the reproductive cycle in sows. We recall that, if the network is consulted for any other type of pig than a sow, the five variables under study should not partake in diagnostic reasoning. If the network is consulted for a sow, on the other hand, any evidence with respect to the reproductive cycle should be taken into consideration upon constructing a diagnosis.

We suppose that the part of the network shown in Figure 1 is consulted for a sow. We recall that the parameter probabilities for the variable *Intra-uterine infection* shown in Table 1 have been conditioned on the pig being a sow and, hence, apply to the current consultation. We observe from these parameter probabilities that both an early and a late intra-uterine infection are more likely to result from an infection with classical swine fever than from a common airways or gastro-intestinal infection. If the sow shows evidence of an intra-uterine infection, therefore, the probability of a viraemia of classical swine fever being present will increase. Also the probabilities of the presence of a common primary infection will increase, yet to a lesser extent. The table further shows that at least one of the various types of infection need be present for the intra-uterine infection to occur. Given evidence of an intra-uterine infection, therefore, a dependence has arisen between the presence of a viraemia of classical swine fever and the presence of a primary other infection. If we would find evidence that the sow does not suffer from any common primary infection, for example, then the probability of a viraemia of classical swine fever being present would increase. We used our Bayesian-network tool to empirically support these considerations. After entering the evidence *Pig type* $=$ *sow* and *Intra-uterine infection* $=$ *early*, we found that the probability of a viraemia of classical swine fever being present equals 0.297. Subsequently entering the evidence *Primary other infection* $=$ *none* served to increase this probability to 1.00, as expected. We would like to note that since in practice it is very unlikely that the presence of an early intra-uterine infection and the absence of any common primary infection are observed with certainty, our example should be taken as serving illustrative purposes only.

Now suppose that the part of the network from Figure 1 is consulted for a boar. Upon entering the type of pig into the network, the five variables capturing issues of reproductivity will be rendered not applicable. We observe that the parameter probabilities for the value *not applicable* of the variable *Intra-uterine infection* are equal to 1 for all combinations of values for the variables *Primary other infection* and *CSF viraemia*. The indirect evidence of the value *not applicable*, therefore, does not allow for distinguishing between these combinations of values and, hence, does not have any effect on the probabilities that are computed for the variables *Primary other infection* and *CSF viraemia*. If we would now find evidence that the boar does not suffer from any common primary infection, for example, then the probability of a viraemia of classical swine fever being present would not change. Once again, we used our Bayesian-network tool to empirically support these considerations. After entering the evidence *Pig type* $=$ *boar*, we found that the probability of a viraemia of classical swine fever being present equals 0.0001. Subsequently entering the evidence *Primary other infection* $=$ *none* indeed did not change this probability, thereby demonstrating that the indirect evidence of the value *not applicable* for the variable *Intra-uterine infection* did not induce a dependence between the variables *Primary other infection* and *CSF viraemia*. Note that evidence for, for example, the variable *Body temperature* would induce such a dependence.

4.3 Summary

To summarise, we have addressed in this section the issue of modelling diverging knowledge for different subpopulations that have to be captured in a single Bayesian network. More specifically, we have focused on knowledge that pertains to just a part of the overall population. To provide for correctly modelling this knowledge, we enhanced the meanings of the variables involved by including a meta-level value capturing applicability. By including this meta-level value, we forestalled the necessity of providing each variable involved with an additional incoming arc in the network's structure to distinguish the subpopulation of in-

terest from other individuals. By introducing the meta-level value, we increased the probability tables for the variables involved in size. However, since the added knowledge does not involve any uncertainty, the additionally required probabilities equal either 0 or 1 and do not need any further assessment by domain experts. By assuring that the parameters probabilities for the meta-level value of a variable are equal for all combinations of values for the variable's possible causes, the meta-level value is guaranteed not to interfere with any reasoning process.

5 Discussion

In close collaboration with two experts, we are building a Bayesian network for the early detection of classical swine fever in pigs. Although our network will ultimately be tailored to detecting the disease in pig herds, we decided to begin with modelling the onset and progression of the disease in individual animals. Having modelled the patterns of classical swine fever that are typically observed in individual pigs, we now are in the process of extending our network to apply to herds. We argued that, as a first step, we had to render our network applicable to the various types of pig distinguished in the domain. We presented a method for including diverging knowledge about these different types in our network. We feel that our method is generally applicable and can be used in any problem domain in which such diverging knowledge has to be captured.

The projected extension of our network raises the question whether we have to develop different networks for the various levels of scope or a single network in which knowledge from the different levels are combined. Little practical experience is available with either of the two approaches, however. To the best of our knowledge, for example, just a few Bayesian networks combine information of several individuals to obtain results about the entire population [7, 8]. For the near future, we envision studying the modelling solutions used in these networks as well as the possibilities and restrictions of using relational models. We feel that the issue of combining knowledge from different levels of scope deserves attention, not just for further developing our network, but also for applications in other domains.

Acknowledgements

This research was (partly) supported by the Netherlands Organisation for Scientific Research (NWO). We are most grateful to Willie Loeffen and Armin Elbers from the Central Institute for Animal Disease Control who spent much effort in the construction of our network for classical swine fever. We would further like to thank Ad Feelders for the fruitful discussion we had about relational networks.

References

[1] P.J.F. Lucas, L.C. van der Gaag, A. Abu-Hanna (2004). Editorial: Bayesian models in biomedicine and health-care. *Artificial Intelligence in Medicine*, vol. 30, pp. 201 – 214.

[2] I.J. McKendrick, G. Gettinby, Y. Gua, S.W.J. Reid, C.W. Revie (2000). Using a Bayesian belief network to aid differential diagnosis of tropical bovine diseases. *Preventive Veterinary Medicine*, vol. 47, pp. 141–156.

[3] L. Otto, C.S. Kristensen (2004). A biological network describing infection with *Mycoplasma hyopneumoniae* in swine herds. *Preventive Veterinary Medicine*, vol. 66, pp. 141–161.

[4] A. Stegeman, A.R.W. Elbers, J. Smak, M.C.M. De Jong (1999). Quantification of the transmission of classical swine fever virus between herds during the 1997-1998 epidemic in the Netherlands. *Preventive Veterinary Medicine*, vol. 42, pp. 219–234.

[5] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B. Aleman, B.G. Taal (1999). How to elicit many probabilities. *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, pp. 647–654.

[6] N. Friedman, L. Getoor, D. Koller, A. Pfeffer (1999). Learning probabilistic relational models. *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann, San Francisco, pp. 1300–1309.

[7] G.F. Cooper, D.H. Dash, J.D. Levander, W-K. Wong, W.R. Hogan, M.M. Wagner (2004). Bayesian biosurveillance of disease outbreaks. *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, AUAI Press, Arlington, pp. 99–103.

[8] E. Jørgensen, N. Toft (1999). A Bayesian network based monitoring sytem for sow management. Dina Research Report No. 80, pp 22.