Discrimination and its sensitivity in probabilistic networks

Silja Renooij and Linda C. van der Gaag
Department of Information and Computing Sciences, Utrecht University
P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
{silja,linda}@cs.uu.nl

Abstract

A probabilistic network built for an application domain often has a single output variable of interest, for which either the posterior probability of one of its values or its most likely value is reported and used for subsequent decision making. For our domain of application, however, we are interested primarily in how well the network distinguishes between various compound output values of interest for different diagnostic variables. To capture this, we introduce a concept of discrimination, and illustrate a measure to this end, based upon joint posterior probabilities. In addition, we address the sensitivity of discrimination to inaccuracies in a network's parameters and show that standard sensitivity functions suffice for studying the effects of such inaccuracies.

1 Introduction

A probabilistic network designed for diagnostic support in an application domain often has a single output variable of interest, capturing the possible diagnoses. Our application domain of classical swine fever, however, aims at multiple-disorder diagnosis. To this end, we have two output variables of interest: a main diagnostic variable to detect classical swine fever, and a secondary variable to capture primary other infections. Although outbreaks of classical swine fever occur seldomly, it is a very serious infectious disease which warrants early detection to prevent rapid spreading. Early detection, however, is hampered by close resemblance of the early symptoms of the disease to those of common infections, and by the simultaneous presence of such infections. A model for early detection of classical swine fever, therefore, needs to be able to distinguish between classical swine fever in an early stage and a primary other infection. Moreover, it should be capable of diagnosing classical swine fever in combination with common infections.

In order to determine how well a probabilistic network can distinguish between different diagnoses in an individual case, it does not always suffice to consider the most likely value of a variable of interest, or its posterior distribution, especially when more than one diagnostic variable is concerned. Therefore, we introduce the concept of *evidence-specific discrimination* between values of one or more variables. Various measures involving posterior probabilities for the diagnoses of interest can be used to capture such discrimination. In this paper we illustrate the concept of discrimination by defining the absolute difference between posterior probabilities as a simple discrimination measure: the further these probabilities are apart, the better the network discriminates between the associated diagnoses.

We note that the term discrimination is somewhat overloaded; it is, for example, often used in the context of classification problems: "can our model discriminate between pigs that have classical swine fever and pigs that have not?" This question, although relevant, concerns discrimination between *cases*, and not between diagnoses in an individual case, which is the problem we address here.

Posterior probabilities can be highly sensitive to changes in a probabilistic network's numerical parameters (Van der Gaag & Renooij, 2001). As the parameters are generally estimated from (incomplete) data or assessed by human experts in the domain of application, they are inevitably inaccurate. To study the robustness of discrimination to parameter inaccuracies, we can study the sensitivity of the output probabilities involved to parameter changes by means of a sensitivity analysis. To

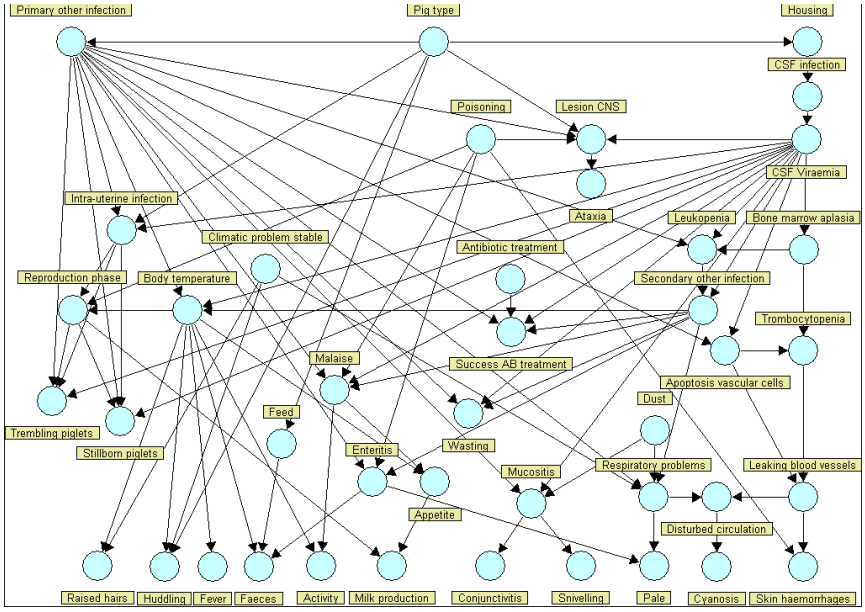

Figure 1: The network for early detection of classical swine fever (csf).

this end, we show how to derive a function that captures the *sensitivity of discrimination* to parameter changes. In addition, we demonstrate that we can efficiently compute a sensitivity function for joint posterior probabilities, required in order to study the dynamics, and therefore robustness, of discrimination between values of two or more variables.

The paper is organised as follows. In Section 2 we describe an application which motivates the need for a concept of discrimination and introduce a preliminary measure to this end. In Section 3, we establish functions that allow for studying the robustness of discrimination to parameter inaccuracies. The paper ends with conclusions and directions for further research in Section 4.

2 The Concept of Discrimination

Classical swine fever (csf) is a serious infectious disease and, although outbreaks occur seldomly, its rapid spreading warrants early detection. Classical swine fever is hard to diagnose in an early stage, due to the high variation in its associated clinical patterns, which strongly resemble those of common other infectious diseases. To support the early detection of csf, a probabilistic network is being developed (Geenen, Elbers, Van der Gaag & Loeffen, 2006). The network, shown in Figure 1, currently includes 82 directed edges, 2113 conditional probabilities, and 41 variables of which 24 can be observed upon clinical investigation. The variables capture processes in the underlying pathogenesis, risk factors, relevant clinical signs, and alternative explanations for these signs.

The main diagnostic variable CSF Viraemia in the network models the presence or absence of csf in an individual pig. The extremely low prior for the presence of csf (0.0000016), in combination with the common occurrence of other infections resembling csf, both in pigs with and without csf, makes that these diseases cannot all be modelled in a single variable: csf would never be diagnosed. A secondary diagnostic variable in the network therefore models primary other infections as possible alternative explanations of a pig's symptoms. As a result, for a given pig, not only the network's prediction of the probability of csf is of interest, but our primary interest is to know how well the network distinguishes csf in an early stage, with or without another infection being present, from just a primary other infection.

Known concepts as the most likely value of a variable of interest, or its posterior distribution, do not always suffice to determine how well a probabilistic network can distinguish between different diagnoses in an individual case, especially when more than one diagnostic variable is concerned. To

capture this, we therefore introduce the novel concept of *evidence-specific discrimination* between two combinations of values for one or more variables. To measure discrimination, we use a function of the posterior probabilities of the (compound) values of interest. This *measure of discrimination* preferably has the property of obtaining a maximum value when one of the posteriors equals zero and the other 1: it is then easy to discriminate between the two associated diagnoses; likewise, the measure should obtain a minimum value when the posteriors are equal. Possible measures of discrimination can be based on (odds) ratios, or more complex functions. In this paper, for the purpose of illustration, we use a simple and straightforward measure of discrimination, defined below. In the remainder of this paper, we will write $\Pr(a \mid e)$, to denote an output probability under study, where a is a specific value assignment to one or more variables A of interest and e denotes the available evidence.

Definition 1. Let $\Pr(a \mid e)$ and $\Pr(b \mid e)$ be two output probabilities of interest. The amount of *discrimination* of the network between a and b in the context of evidence e, written $d(a\,; b \mid e)$, equals $\mid \Pr(a \,|\, e) - \Pr(b \,|\, e) \mid$.

The above measure $d(a\,; b \mid e)$ takes on values between zero and 1, with larger values indicating a larger amount of discrimination. This specific measure also has the benefit of being symmetric in its arguments, that is, $d(a; b \,|\, e) = d(b; a \,|\, e)$.

Example 1. Discrimination can be measured between different values of the same variable. For example, discrimination between a gastro-intestinal infection (value gi of variable POI, modelling primary other infections) and an airway infection (value ai of variable POI) in pig 14 is given by

$$\begin{aligned} d(gi; ai \,|\, 14) &= \mid \Pr(gi \,|\, 14) - \Pr(ai \,|\, 14) \mid \\ &= 0.54 - 0.10 = 0.44 \end{aligned}$$

indicating that the network can easily distinguish between these two types of infection in this pig. Discrimination can also be studied for values of different variables. For example, discrimination between classical swine fever (value csf of variable CSF) and a gastro-intestinal infection in pig 169 is given by $d(csf; gi \,|\, 169) = \mid \Pr(csf \,|\, 169) - \Pr(gi \mid 169) \mid = 0.20 - 0.13 = 0.07$, indicating that the network has some difficulty distinguishing csf from a common infection in this pig. Discrimination can even be studied for value assignments to more than one variable. For example, discrimination between the presence and absence of classical swine fever in combination with an airway infection in pig 304: $d(csf, ai\,; \neg csf, ai \mid 304) = \mid \Pr(csf, ai \mid 304) - \Pr(\neg csf, ai \,|\, 304) \mid = \mid 0.01 - 0.15 \mid = 0.14$; this indicates that the network is capable of diagnosing csf in combination with another infection in this pig. □

3 Robustness of Discrimination

Sensitivity analysis is a powerful tool for studying the robustness of a probabilistic network's output probabilities to inaccuracies in the network parameters. Since discrimination is defined in terms of output probabilities, its robustness to parameter changes is a relevant matter, and can be studied by means of the functions that result from a sensitivity analysis. We now review some known properties of such sensitivity functions.

3.1 Sensitivity Functions

Sensitivity analysis of a probabilistic network amounts to establishing, for each of the network's numerical parameters, the *sensitivity function* that expresses an output probability of interest in terms of that parameter. Let $x = p(b \mid \pi)$ be a parameter under study, where b is a value of some variable B and π is a combination of values for B's parents. We now use $f_a^e(x)$ to denote the sensitivity function that expresses the output probability $\Pr(a \mid e)$ in terms of the parameter x.

Any one-way sensitivity function $f_a^e(x)$ is a quotient of two linear functions in the parameter x under study (Castillo, Gutiérrez & Hadi, 1997; Coupé & Van der Gaag, 2002). More formally, the function takes the form

$$f_a^e(x) = \frac{\Pr(a,\, e)(x)}{\Pr(e)(x)} = \frac{c_1 \cdot x + c_2}{c_3 \cdot x + c_4}$$

where the constants c_j, $j = 1, \ldots, 4$, are built from the assessments for the parameters that are not being varied[1]. Efficient algorithms exist to compute

[1] We assume that the parameters pertaining to the same conditional distribution as the parameter under study are co-varied proportionally (Kjærulff & Van der Gaag, 2000).

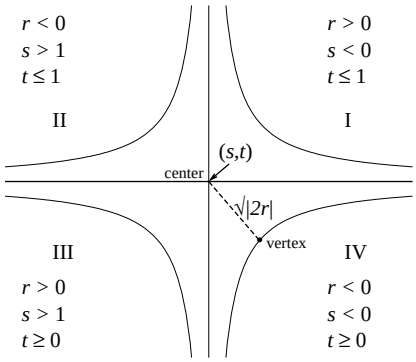

Figure 2: Two hyperbolas with their branches and associated constants (the constraints on s and t are specific for sensitivity functions).

the constants of any sensitivity function relating a (posterior) probability for a value of a single output variable to a network parameter (Coupé & Van der Gaag, 2002; Kjærulff & Van der Gaag, 2000).

The sensitivity function $f_a^e(x)$ can take one of three general forms. The function is *linear* for prior probabilities of interest, or if $\Pr(e)$ is unaffected by the parameter variation ($c_3 = 0$); if $c_4 = 0$, then $c_2 = 0$ and the function reduces to a *constant*. In all other cases the function is a fragment of a *rectangular hyperbola*, which takes the general form

$$f(x) = \frac{r}{x - s} + t$$

where, for a sensitivity function with $c_1, \ldots, c_4$ as before, $s = -c_4/c_3$, $t = c_1/c_3$, and $r = (c_2/c_3) + s \cdot t$. In the remainder of the paper, we assume any sensitivity function to be hyperbolic.

Figure 2 illustrates that a rectangular hyperbola in general has two branches, and two asymptotes defining its center (s, t). We observe that a sensitivity function is defined by $0 \leq x, f(x) \leq 1$; the two-dimensional space of feasible points thus defined, is termed the *unit window*. Since a sensitivity function moreover should be continuous for $x \in [0, 1]$, its vertical asymptote necessarily lies outside the unit window. A hyperbolic sensitivity function therefore is a fragment of a single hyperbola branch.

3.2 Discrimination Dynamics: Simple Values

The robustness of a network's discrimination between a and b, in the context of evidence e, to changes in a parameter x, can be captured by considering $d(a; b \,|\, e)$ as a function of x. In this section we assume that a and b are *simple* values, that is values for a single variable A and a single variable B; the case where a and b are *compound* values is considered in Section 3.3. In this paper we assume that $d(a; b \mid e)$ itself is a function involving simple operators as the sum, the difference, and/or the ratio of posterior probabilities. We will demonstrate, for our choice of measure, that *discrimination sensitivity* $d(a; b|e)(x)$ can now again be described in terms of a rectangular hyperbola.

Proposition 1. *Let* $f_a^e(x) = r_a/(x - s) + t_a$ *and* $f_b^e(x) = r_b/(x-s)+t_b$ *be two sensitivity functions. Then*

$$f_a^e(x) - f_b^e(x) = \frac{(r_a - r_b)}{x - s} + (t_a - t_b)$$

The above immediately follows from having the same constant s in both sensitivity functions, which is justified by the following lemma.

Lemma 1. *For a fixed parameter* x *and evidence* e*, all sensitivity functions* $f_A^e(x)$ *for any variable* A *have the same vertical asymptote.*

Recall that the constant s for a sensitivity function $f_a^e(x)$ equals $s = -c_4/c_3$, where $c_3 \cdot x + c_4 = \Pr(e)(x)$. Given a parameter x, constant s therefore relates to just the available evidence and is independent of the output variable of interest.

Although the difference function from Proposition 1 again is a fragment of one hyperbola branch for $x \in [0, 1]$, it will in general not be a sensitivity function since it can be negative on $[0, 1]$; for our choice of measure, $d(a\,; b \mid e)(x)$ is the absolute value of this difference. For a fixed parameter x and evidence e, establishing the constants of all sensitivity functions $f_A^e(x)$ for any single variable A, rather than for one specific A, comes at no additional computational expense. Establishing $d(a; b|e)(x)$ hence requires no additional network propagations.

The function $d(a; b \mid e)(x)$ now details how discrimination is affected by parameter variation. Discrimination is robust to parameter inaccuracies if its change upon varying a parameter is limited. To assess robustness, we now define different intervals of parameter values having different effects on discrimination.

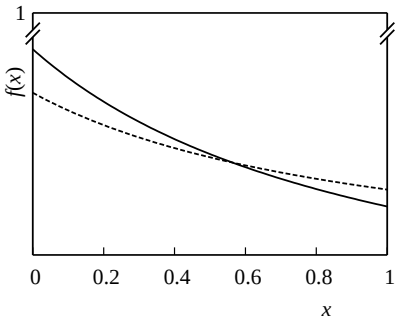

Figure 3: Example sensitivity functions with $s = -1$, $x_{\text{int}} = 0.56$ and $x_{\text{max}} = 0$.

Since $d(a\,;b \mid e)(x)$ is based on two sensitivity functions, which are continuous and monotone for $x \in [0,1]$, we have that maximum discrimination is found on the boundaries of the unit window, that is, for either $x = 0$ or $x = 1$. The value of parameter x where discrimination is maximised will be denoted by x_{max}:

$$x_{\text{max}} = \operatorname*{argmax}_{x \in [0,1]} d(a;b \,|\, e)(x) \in \{0,1\}$$

If the two sensitivity functions $f_a^e(x)$ and $f_b^e(x)$ for the posterior probabilities under consideration intersect within the unit window, such as in Figure 3, then minimum discrimination is attained at this intersection point. Assuming that the two hyperbolas are truly different functions, that is $f_a^e(x) \neq f_b^e(x)$, they intersect for at most one value of x, denoted x_{int}. For our choice of discrimination measure this minimum value equals zero and is attained at

$$d(a;b \,|\, e)(x_{\text{int}}) = 0 \iff x_{\text{int}} = \frac{r_a - r_b}{t_b - t_a} + s$$

If $x_{\text{int}} \in \langle 0,1 \rangle$, then parameter values on opposite sides of x_{int} will result in the same amount of discrimination between the values a and b under consideration (see Figure 4). Let x_{sim} denote the value of x for which discrimination equals the original amount of discrimination between a and b in context e, that is

$$d(a;b \,|\, e)(x_{\text{sim}}) = d(a;b \,|\, e)(x_0),$$

where x_0 is the original value of the parameter x under consideration. For our example measure, the

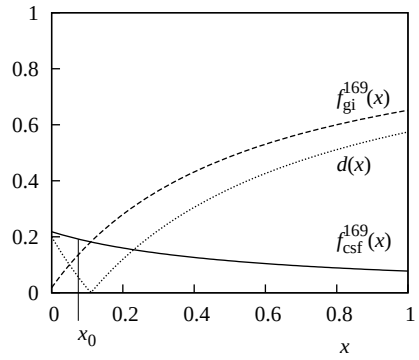

Figure 4: Discrimination $d(csf\,;gi \mid 169)(x)$ for a parameter x with $x_0 = 0.075$; $x_{\text{int}} = 0.11$ and $x_{\text{sim}} = 0.16$.

value of x_{sim} can be easily established from $d_0 = f_a^e(x_0) - f_b^e(x_0)$:

$$x_{\text{sim}} = \frac{r_a - r_b}{t_b - t_a - d_0} + s$$

We now note that x_0 and x_{sim} necessarily lie on opposite sides of x_{int} if the latter two lie within the unit window. As a result, for x-values between x_0 and x_{sim}, discrimination will become smaller than its original value at x_0; we will then say that discrimination decreases, even though it is not a decreasing function of x; for x-values outside the interval bounded by x_0 and x_{sim} discrimination increases. If $x_{\text{sim}} \notin [0,1]$, for example as in Figure 3 with $x_0 = 0.10$ and $x_{\text{sim}} = 1.70$, then, necessarily, x_{max} lies on the same side of x_{int} as x_0, so variation of x from x_0 towards x_{max} increases discrimination, whereas discrimination will become less when x is varied in the opposite direction.

The intersection point of the two hyperbolas does not necessarily lie within the unit window (see for example Figure 5). If the intersection point lies outside the unit window, or if the hyperbola branches do not intersect at all, then discrimination $d(a\,;b \mid e)(x)$ is monotone for $x \in [0,1]$, obtaining its minimum value at $1 - x_{\text{max}}$.

We now have all the ingredients to describe the effect of parameter variation on discrimination.

Proposition 2. *Let $f_a^e(x)$, $f_b^e(x)$, x_{max}, x_{int}, and x_{sim} be as before. Then the network's discrimination $d(a\,;b \mid e)$ between a and b in the context of evidence e changes as follows, upon varying parameter x:*

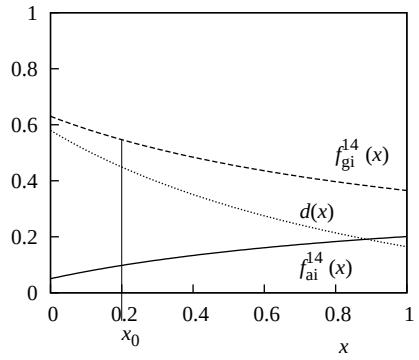

Figure 5: Discrimination $d(gi\,;\, ai \mid 14)(x)$ for a parameter x with $x_0 = 0.20$; $x_{\text{int}} = 2.08$, $x_{\text{sim}} = -1.17$, and $x_{\max} = 0$.

- *if* $x_{\text{int}} \notin [0,1]$ *or* $x_{\text{sim}} \notin [0,1]$*, then discrimination is non-decreasing if* x *is varied towards* $x_{\max}$*, and non-increasing otherwise.*
- *if* $x_{\text{int}} \in [0,1]$ *and* $x_{\text{sim}} \in [0,1]$*, then discrimination is non-decreasing if* x *is varied to* $x \leq \gamma$ *or to* $x \geq \delta$*, where* $[\gamma, \delta] = [\min\{x_0, x_{\text{sim}}\}, \max\{x_0, x_{\text{sim}}\}]$*, and non-increasing otherwise.*

Example 2. Reconsider the network for early detection of csf and its discrimination between csf and a gastro-intestinal infection for pig 169, $d(csf\,;\, gi \mid 169)$. Discrimination between these two values as a function of a parameter x pertaining to the success of treatment with antibiotics ($x_0 = 0.075$), is captured by the following two sensitivity functions:

$$f^{169}_{csf}(x) = \frac{0.12}{x + 0.55} \text{ and } f^{169}_{gi}(x) = \frac{-0.54}{x + 0.55} + 1$$

and shown in Figure 4. From these functions we find that $x_{\text{int}} = 0.11$ and $x_{\text{sim}} = 0.16$. As a result, if a more accurate assessment of the parameter turns out smaller than 0.075, then the network will be able to better discriminate between the two infections in this pig; if a more accurate assessment is larger, but not as large as 0.16, then discrimination becomes worse. Similarly, discrimination $d(gi\,;\, ai \mid 14)$ between a gastro-intestinal and an airway infection in pig 14, as a function of a parameter x pertaining to faeces samples ($x_0 = 0.20$), is captured by the following two sensitivity functions:

$$f^{14}_{gi}(x) = \frac{0.69}{x + 1.19} + 0.05, \text{ and}$$

$$f^{14}_{ai}(x) = \frac{-0.39}{x + 1.19} + 0.38$$

shown in Figure 5. From the formulas we find that $x_{\text{int}} = 2.08$ and $x_{\text{sim}} = -1.17$. We conclude that discrimination decreases for any parameter value larger than 0.20, and increases otherwise. □

For our example measure we have studied the difference between two sensitivity functions, each describing a posterior probability for a simple output value as a function of a network parameter. We note that the difference between two such posterior probabilities in relation to changes in a network parameter has been studied before by Chan & Darwiche (2002) in the context of parameter tuning. They demonstrated that parameter values which enforce a constraint on the difference, or on the ratio, of two posterior probabilities can be computed from partial derivatives established from the network without explicitly determining a sensitivity function. Establishing the constants of the sensitivity function, however, is just as efficient and has the benefit of providing insight in the effects of arbitrary parameter changes on an output of interest.

3.3 Discrimination Dynamics: Compound Values

In this section, we extend the results on the robustness to parameter inaccuracies of discrimination between simple values, to apply to compound value combinations for two or more variables. This type of robustness is relevant in case we need to distinguish between diagnoses for multiple disorders, modelled in separate diagnostic variables. In practice, the number of variables under consideration should typically be small, for computational reasons as well as for interpretability. Although results in this section apply to compound values for any number of variables, we limit the discussion to only two.

We consider the variables A and B and the posterior probability $\Pr(a,\, b \mid e)$ of the compound value ab. A sensitivity function $f^{e}_{ab}(x)$ for ab in the context of evidence e, as a function of x, is readily determined by one of the following two approaches:

I) Extend the network with a new variable Y, with parents A and B and all their compound values as possible values for Y; for the CPT, define $p(y \mid a,\, b) = 1$ iff $y \equiv ab$, and

$p(y \mid a, b) = 0$ otherwise. Enter evidence e into the new network and compute $f^e_{Y=ab}(x)$.

II) Enter evidence e into the original network to compute $f^e_b(x)$, then enter additional evidence b to compute $f^{be}_a(x)$. Finally, multiply the two functions:

$$\begin{aligned} f^e_{ab}(x) &= f^{be}_a(x) \cdot f^e_b(x) \\ &= \left(\frac{r_a}{x - s_{be}} + t_a\right) \cdot \left(\frac{r_b}{x - s_e} + t_b\right) \end{aligned}$$

The first approach requires less propagations, but in a more complex network, and establishes the sensitivity functions for all compound values of the variables under consideration. The second approach leaves the network as-is and provides all information for establishing the sensitivity functions $f^e_{Ab}(x)$ for all values of all variables A in the network. The multiplication step is simplified by the observation that the resulting function is again a sensitivity function and therefore a rectangular hyperbola. This indeed follows after careful inspection of all constants involved.

Proposition 3. *Let $f^{be}_a(x) = r_a/(x - s_{be}) + t_a$ and $f^e_b(x) = r_b/(x - s_e) + t_b$ be two sensitivity functions. Then*

$$f^e_{ab}(x) = \frac{r_a \cdot t_b + (s_e - s_{be}) \cdot t_a \cdot t_b}{x - s_e} + t_a \cdot t_b$$

is the sensitivity function relating the joint probability $\Pr(a, b \mid e)$ to parameter x.

Proof. First we rewrite the formulas for the hyperbolic sensitivity functions in terms of a fraction of linear functions:

$$f^{be}_a(x) = \frac{\Pr(a, b, e)(x)}{\Pr(b, e)(x)} = \frac{c_1 \cdot x + c_2}{c_3 \cdot x + c_4}$$

where $-c_4/c_3 = s_{be}$, $c_1/c_3 = t_a$, and $c_2/c_3 = r_a - s_{be} \cdot t_a$, and

$$f^e_b(x) = \frac{\Pr(b, e)(x)}{\Pr(e)(x)} = \frac{c_3 \cdot x + c_4}{c_5 \cdot x + c_6}$$

where $-c_6/c_5 = s_e$, $c_3/c_5 = t_b$, and $c_4/c_5 = r_b - s_e \cdot t_b$. Now,

$$\begin{aligned} f^e_{ab}(x) &= f^{be}_a(x) \cdot f^e_b(x) = \frac{c_1 \cdot x + c_2}{c_5 \cdot x + c_6} \\ &= \frac{r_{ab}}{x - s_{ab}} + t_{ab} \end{aligned}$$

where

$$\begin{aligned} s_{ab} &= -\frac{c_6}{c_5} = s_e \\ t_{ab} &= \frac{c_1}{c_5} = \frac{c_1}{c_3} \cdot \frac{c_3}{c_5} = t_a \cdot t_b \\ r_{ab} &= \frac{c_2}{c_5} + s_e \cdot t_a \cdot t_b = \frac{c_2}{c_3} \cdot \frac{c_3}{c_5} + s_e \cdot t_a \cdot t_b \\ &= (r_a - s_{be} \cdot t_a) \cdot t_b + s_e \cdot t_a \cdot t_b \end{aligned}$$

□

From the above observations, we have that all properties for sensitivity functions and discrimination derived in the previous section readily apply to the compound values case.

Example 3. Reconsider the network for early detection of csf. We study the network's discrimination, and its robustness, between csf and one of the primary other infections. More specifically, since other infections are quite common in pigs, we are interested in whether or not csf can be distinguished from them. In this example, we focus on the difference between $\Pr(csf, \neg gi \mid 169)$ and $\Pr(\neg csf, gi \mid 169)$ for pig 169. The robustness of discrimination, as a function of a parameter x, can be studied by means of the corresponding sensitivity functions: $|f^{169}_{csf, \neg gi}(x) - f^{169}_{\neg csf, gi}(x)|$. The constants for the rectangular hyperbola representing this difference are found by applying Propositions 3 and 1 to the functions $f^{169, \neg gi}_{csf}(x)$, $f^{169, gi}_{\neg csf}(x)$, and $f^{169}_{gi}(x)$ and exploiting the fact that $f^{169}_{\neg gi}(x) = 1 - f^{169}_{gi}(x)$. From these constants, x_{int} and x_{sim} can be straightforwardly computed.

Examples of the sensitivity functions for the simple output values under consideration and a parameter x pertaining to the success of an antibiotics treatment are given in Figure 6. The sensitivity functions for the compound values of interest for the same x, together with discrimination as a function of x, are shown in Figure 7. Note that Figure 7 gives valuable insight into the dynamics of discrimination between csf and gastro-intestinal infections, which is not obvious from Figure 6: although from Figure 6 we can see that changes in the posterior probability of gi will pull the probabilities for its combination with csf towards the center of the probability range, it is not immediately obvious from this figure that the functions for the compound values will intersect, nor where this will occur. □

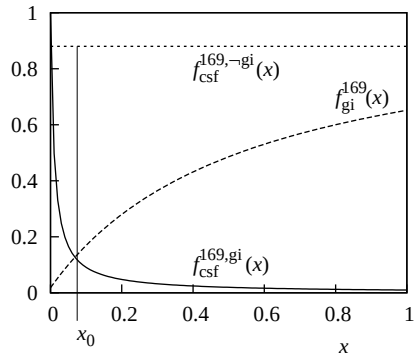

Figure 6: Sensitivity functions for simple values of variables CSF and POI, for a parameter x pertaining to the success of an antibiotics treatment.

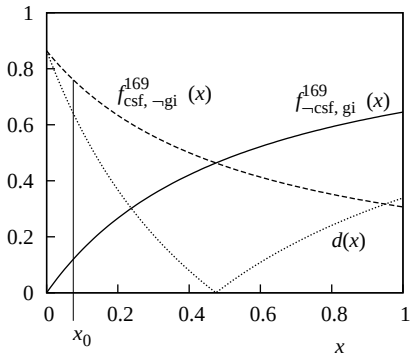

Figure 7: Sensitivity functions for compound values of variables CSF and POI, for a parameter x pertaining to the success of an antibiotics treatment, together with discrimination $d(csf, \neg gi; \neg csf, gi \mid 169)(x)$ for parameter x.

4 Conclusions

In this paper we introduced the concept of evidence-specific discrimination to investigate how well a probabilistic network can distinguish between two or more output values or, more in general, between value combinations for two or more output variables of interest. We illustrated a simple measure of discrimination based on the difference between two posterior probabilities. Subsequently, we demonstrated how sensitivity functions can be employed to study the robustness of discrimination to parameter inaccuracies, even when discrimination concerns compound values rather than simple ones.

Our results on the dynamics of discrimination build to a large extent on the observation that, in the same evidence context, simple operations on hyperbolic sensitivity functions for the same parameter x, again result in a rectangular hyperbola. This entails that more sophisticated discrimination measures, such as for example (odds) ratios or $|f_a^e(x) - f_b^e(x)|/(f_a^e(x) + f_b^e(x))$, can be straightforwardly employed with the techniques presented in this paper. Further research is required to investigate what measure of discrimination is most suitable in what situation, and what amount of discrimination is acceptable or desirable. In addition, we plan on investigating to what extent results that address evidence-dependent bounds on sensitivity functions (Renooij & Van der Gaag, 2005) can be employed to make general statements concerning discrimination involving all network parameters.

References

E. Castillo, J.M. Gutiérrez, A.S. Hadi (1997). Sensitivity analysis in discrete Bayesian networks. *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 27, pp. 412 – 423.

H. Chan, A. Darwiche (2002). When do numbers really matter? *Journal of Artificial Intelligence Research*, vol. 17, pp. 265 – 287.

V.M.H. Coupé, L.C. van der Gaag (2002). Properties of sensitivity analysis of Bayesian belief networks. *Annals of Mathematics and Artificial Intelligence*, vol. 36, pp. 323 – 356.

P.L. Geenen, A.R.W. Elbers, L.C. van der Gaag, W.L.A. Loeffen. Development of a probabilistic network for clinical detection of classical swine fever. *Proceedings of the 11th Symposium of the International Society for Veterinary Epidemiology and Economics*, Cairns, Australia, pp. 667669, 2006.

U. Kjærulff, L.C. van der Gaag (2000). Making sensitivity analysis computationally efficient. *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, pp. 317 – 325.

L.C. van der Gaag, S. Renooij (2001). Analysing sensitivity data from probabilistic networks. *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann, San Francisco, pp. 530 – 537.

S. Renooij, L.C. van der Gaag (2005). Exploiting evidence-dependent sensitivity bounds. *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence*, AUAI Press, Corvallis, OR, pp. 485-492.