

Diagnosing Lyme disease - Tailoring patient specific Bayesian networks for temporal reasoning

Kristian G. Olesen (kgo@cs.aau.dk)

Department of Computer Science, Aalborg University
Fredrik Bajersvej 7 E, DK-9220 Aalborg East, Denmark

Ole K. Hejlesen (okh@mi.aau.dk)

Department of Medical Informatics, Aalborg University
Fredrik Bajersvej 7 D, DK-9220 Aalborg East, Denmark

Ram Dessau (rde@cn.stam.dk)

Clinical Microbiology, Næstved Hospital
Denmark

Ivan Beltoft (ibe@netcompany.com) and Michael Trangeled (michael@trangeled.com)

Department of Medical Informatics, Aalborg University

Abstract

Lyme disease is an infection evolving in three stages. Lyme disease is characterised by a number of symptoms whose manifestations evolve over time. In order to correctly classify the disease it is important to include the clinical history of the patient. Consultations are typically scattered at non-equidistant points in time and the probability of observing symptoms depend on the time since the disease was inflicted on the patient.

A simple model of the evolution of symptoms over time forms the basis of a dynamically tailored model that describes a specific patient. The time of infliction of the disease is estimated by a model search that identifies the most probable model for the patient given the pattern of symptom manifestations over time.

1 Introduction

Lyme disease, or Lyme Borreliosis, is the most commonly reported tick-borne infection in Europe and North America. Lyme disease was named in 1977 when a cluster of children in and around Lyme, Connecticut, displayed a similar disease pattern. Subsequent studies revealed that the children were infected by a bacteria, which was named *Borrelia burgdorferi*. These bacteria are transmitted to humans by the bite of infected ticks (*Ixodus* species). The clinical manifestations of Lyme borreliosis are described in three stages and may affect different organ systems, most frequently the skin and nervous system. The diagnosis is based on the clinical findings and results of laboratory testing for an-

tibodies (immunoglobulin M and G) in the blood. Methods for direct detection of the bacteria in the tissue or the blood are not available for routine diagnosis. Thus it may be difficult to diagnose Lyme disease as the symptoms may not be exclusive for the disease and testing for antibodies may yield false positive or negative results.

In the next section we give an overview of Lyme disease and in section 3 a brief summary of two earlier systems are given. In section 4 we suggest an approach that overcome the difficulties by modelling multiple consultations spread out in non-equidistant points in time. This approach assumes that the time of infection is known, but this is rarely the case. We aim for an estimation of this fact; a number of hypothesised models are generated and the best one is

Organ system	Stage 1: Incubation period one week (few days to one month)	Stage 2: Incubation period one week to a few months	Stage 3: Incubation period months to years
Skin	Erythema migrans	Borrelial lymphocytoma	Acrodermatitis chronica atrophicans
Central nervous system		Early neuroborreliosis	Chronic neuroborreliosis
Joints			Lyme arthritis
Heart		Lyme carditis	
Average sensitivity of antibody detection (IgG or IgM)	50%	80%	100%

Table 1: Clinical manifestations of Lyme disease.

identified based on the available evidence. In section 5 a preliminary evaluation of the approach is described and finally we conclude in section 6.

2 Lyme Disease

According to European case definitions (http://www.oeghmp.at/eucalb/diagnosis_case-definition-outline.html) Lyme disease and its manifestations are described in three stages as shown in Table 1.

Erythema migrans is the most frequent manifestation. The rash which migrates from the site of the tickbite continuously grows in size. Erythema migrans is fairly characteristic and the diagnosis is clinical. Laboratory testing is useless as only half of the patients are positive, when the disease is localized to the skin only. Neuroborreliosis is less frequent, in Denmark the incidence is around 3/100.000 per year. The IgM or IgG may be positive in 80% of patients with neuroborreliosis, but if the duration of the the clinical disease is more than two months, then 100% of patients are positive. Figure 1 shows that not only the probability of being positive, but also the level of antibodies vary as a function of time. Thus the laboratory measurement of IgM and IgG antibodies depends both on the duration of the clinical disease and the dissemination of the infection to other organ systems than the skin. The incidence of the

other manifestations of Lyme disease is more rare. The diagnosis of lyme arthritis is especially difficult as the symptoms are similar to arthritis due to other causes and because the disease is rare.




Figure 1: IgM and IgG development after infection of Lyme disease.

The epidemiology is complex. For example different age groups are at different risks, there is a large seasonal variation due to variations in tick activity (Figure 2), the incubation period may vary from a few days to several years and some clinical manifestations are very rare. There are large variations in incubation time and clinical progression. Some patients may have the disease starting with a rash (erythema migrans) and then progressing to neuroborreliosis, but most patients with neuroborreliosis are not aware of a preceding rash. The description of the disease and its progression primarily based on (Smith et al., 2003) and (Gray et al., 2002).

There are also large individual variations in

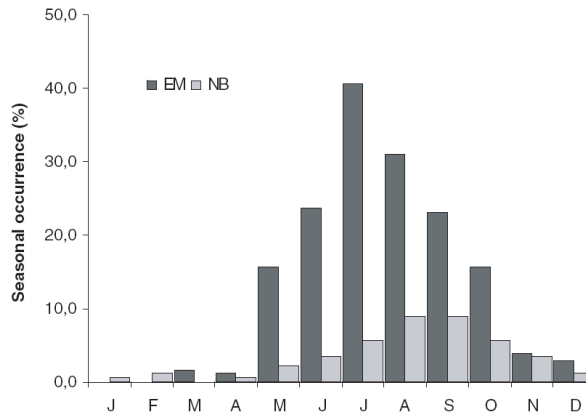


Figure 2: Comparison of seasonal occurrence of erythema migrans (EM, $n = 566$) and neuroborreliosis (NB, $n = 148$) in northeastern Austria, 1994-1995. (Gray et al., 2002).

the behaviour of the patient and the physician, about when to visit the doctor, when to take the laboratory test ect. Possible repeated visits to the doctor and repeated laboratory testing are performed at variable intervals. Joint pain and arthritis is a very common problem in the general population, but is only very rarely due to Lyme borreliosis. However, many patients with arthritis are tested for Lyme borreliosis (30% of all patients samples) and as the prior probability of Lyme disease is low, the posterior probability of Lyme arthritis is also low in spite of a positive test result. Laboratories in some countries attempt to improve the diagnosis of Lyme disease by using a two-step strategy in antibody detection, by setting a low cut-off on the screening test and then performing a more expensive and complex test to confirm or discard the result of the primary test. As the second test is using the same principle of indirect testing for antibody reactions to different Borrelial antigens the two tests are highly correlated, and the information gain is limited. Thus, the basic problem of false positive or false negative results is not solved. It was therefore found important to develop a decision support system to assist the clinician by calculating the posterior probability of Lyme disease to guide the choice of treatment. An evidence based clinical diag-

nosis is supported by incorporating clinical and laboratory data into the model.

To capture the complex patterns of the disease, a model must incorporate the relevant clinical evidence including estimation of the temporal aspects of time since the tickbite, the duration of clinical disease and the development of antibody response.

3 Existing Models for Diagnosis of Lyme Disease

We have knowledge of two models that have been developed to assist the medical practitioner in the diagnosis of Lyme disease (Dessau and Andersen, 2001; Fisker et al., 2002). Both models are based on Bayesian networks (Jensen, 2001) and as the latter is a further development of the former, they share most of the variables.

The models include a group of variables describing general information and knowledge including age and gender of the patient, the patient's exposure to ticks (is the patient a forrest worker or orienteer), the month of the year and whether the patient recalls a tick bite. The information is used to establish whether or not the conditions for a Lyme disease infection has been present. This part of the model influence the hypothesis variable, Borrelia.

Another section of the models describe clinical manifestations. The findings are influenced by the hypothesis variable, but may be caused by other reasons. Similarly, a section of variables describing laboratory findings may be caused by either Lyme disease or by other disorders.

The structure of the model by Dessau and Andersen (2001) is shown in Figure 3. In this model the three stages of Lyme disease is explicitly represented as three copies of the hypothesis and the findings sections. The conditional probabilities reflect the temporal evolution of symptoms, such that e.g. neuroborreliosis is more probable in stage two than in stages one and three. A problem with this approach is that the progression of the disease is uncertain, and consequently it is unclear which copy of e.g. finding

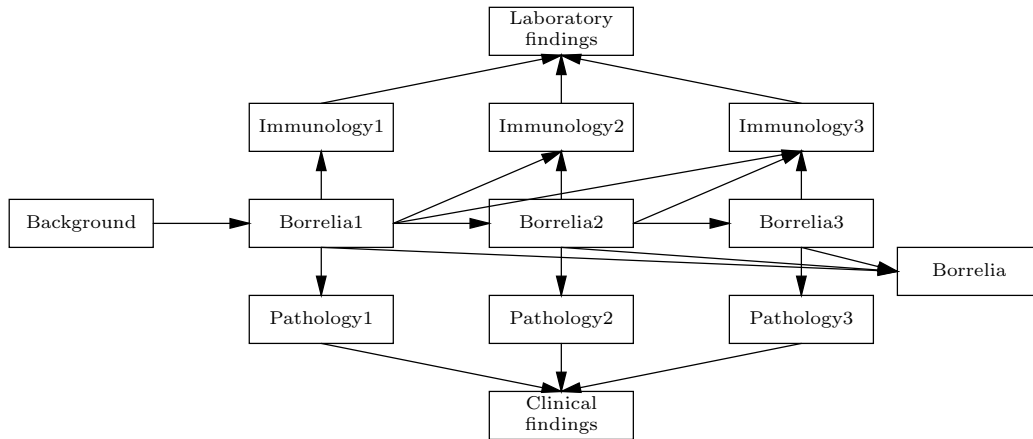


Figure 3: Structure of the Bayesian network by Dessau and Andersen.

variables to use. This was resolved by combining variables from each stage in single variables that describe the observable clinical manifestations and laboratory findings.

The progress of Lyme disease was modelled by a variable “duration” (not shown in the figure), indicating the time a patient has experienced the symptoms. It does not distinguish between the duration of the different symptoms, but models the total duration of illness. If the duration is short it indicates a stage one Lyme disease and so forth.

Dessau and Andersen’s model is a snapshot of a patient at a specific point in time and it does not involve temporal aspects such as means for entering of multiple evidence on the same variables as a result of repeated consultations.

The model by Fisker et al. (2002) aims to include these issues. The idea is to reflect the clinical practice to examine a patient more than once if the diagnosis is uncertain. This gives the medical practitioner the opportunity to observe if the development of the disease corresponds to the typical pattern. The model mainly includes the same variables, but the structure of the network is different (see Figure 4).

Instead of triplicating the pathology for the different stages of the disease the temporal aspect was incorporated in the model by defining the states of manifestation variables as time intervals. This approach was inspired by (Arroyo-Figueroa and Sucar, 1999) that intro-

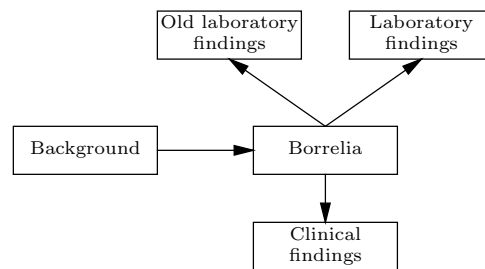


Figure 4: Structure of the Bayesian network by Fisker et al.

duced Temporal Nodes Bayesian Networks as an alternative to dynamic Bayesian networks. The time intervals represent the time since the symptom was first observed, and both the length and the number of time intervals was tailored individually for each node. With this approach it became possible to enter “old” evidence into the model. For example, if erythema migrans was observed on, or recalled by, the patient ten weeks ago and lymphocytoma is a current observation, the *EM* variable is instantiated to the state *8 weeks - 4 months* and the *lymphocy* is instantiated to *< 6 weeks*.

The model also incorporated the ability to enter duplicate results of laboratory tests. This option was included by repeating the laboratory findings and including a variable specifying the time between the result of the old test and the current test (not shown in the figure).

The two models basically use the same features but differ in the modelling of the progres-




Figure 5: The structure of the tailored patient specific model is composed of background knowledge and consultation modules.

sion of the disease. Further, the model by Fisker et al. takes the use of data from previous consultations into consideration.

Both models implement the possibility of reading the probability of having Lyme disease in one of the three stages, but with two different approaches. The model by Fisker et al. models the stages as states in a single variable, whereas the model by Dessau and Andersen models the stages explicitly as causally dependent variables.

The main problem in both models is the temporal progression of Lyme disease. Dessau and Andersen’s model gives a static picture of the current situation, that only takes the total duration of the period of illness into consideration. Fisker et al. is more explicit in the modelling of the progress of the disease by letting states denote time intervals since a symptom was first observed and by duplicating variables for laboratory findings. This enables inclusion of findings from the past, but is still limited to single observations for clinical evidence and two entries for laboratory tests.

Besides the difficulties in determining the structure of the models, a considerable effort was required to specify the conditional probabilities.

In the following section we propose a further elaboration of the modelling by introduction of continuous time and an arbitrary number of consultations.

4 Tailoring patient specific models

We aim for a Bayesian network that incorporates the full clinical history of individual patients. The model is composed of mod-

ules, where each module describes a consultation. The problem is that consultations may appear at random points in time and that the conditional probabilities for the symptoms at each consultation vary, depending on the time that has passed since the tick bite. Therefore frameworks such as dynamic Bayesian networks and Markov chains do not apply. We introduce continuous time inspired by (Nodelman et al., 2002). The proposed approach involve a model for the conditional probabilities, but before reaching that point we determine the structure of the model.

4.1 Structure

The structure of the patient specific model is illustrated in Figure 5. At the first consultation the background knowledge module is linked to a generic consultation module consisting of the disease node, clinical findings and laboratory findings. In order to keep focus on the overall modelling technique we deliberately kept the model simple. The consultation module is a naive Bayes model, and subsequent consultations are included in the the model by extendeding it with a new consultation module. This process can be repeated for an arbitrary number of consultations. The consultation modules are connected only through the disease nodes. This is a quite strong assumption that may be debated, but it simplifies things and keep the complexity of the resulting model linear in the number of consultations. Less critical is that the disease nodes do not take the stage of the disease into account; we consider that the classification into stages are mostly for descriptive purposes, but it could be modeled as stages in

the disease node, although this would complicate the quantitative specification slightly.

4.2 Conditional probability tables

As the symptoms of Lyme disease vary over time we assume a simple model for the conditional probabilities in the model. We approximate the conditional probabilities for the symptoms by a continuous function, composed as a mixture of two sigmoid functions. This is a somewhat arbitrary choice; other models may be investigated, but for the present purpose this simple model suffice. The functions describing the conditional probabilities are based on empirical knowledge. From existing databases we extract the probabilities for the various symptoms at different times since the infection was inflicted on the patient. The time of infliction is usually not known, but around one third of the patients recall a tick bite. In other cases there is indirect evidence for the time of the bite, such as limited periods of exposure.

An example is shown in Figure 6, where the resulting functions for the probability of showing the symptom EM up to 100 days after the infection are drawn. As can be seen, the age of the patient has been incorporated as a parameter to the function. Thus, we can compute the conditional probabilities for a consultation module directly, provided that the time since the infection is known. This is rarely the case.

4.3 Determining the time of infection

The purpose of the Bayesian model is to calculate the probability of Lyme disease based on the clinical findings and possible laboratory evidence. As described in the previous section, the conditional probabilities of the symptoms and the serology are determined by the time since the tick bite. Thus, in order to construct the complete model it is necessary to know the time of infliction.

The clinical history for a given patient is illustrated in Figure 7. Different time intervals are shown in the figure. The interval t_1 is the time from the infection to the first consultation, t_2 is the interval between the first and the second consultation and so on. When the pa-




Figure 6: The probability of experiencing erythema migrans (EM) over time modelled as continuous functions. The development of EM is dependent on the age of the patient.




Figure 7: Time line that represents the consultations that forms the clinical history.

tient consults the medical practitioner the time of the consultation is registered. Therefore, the intervals between the consultations are known, whereas the interval, t_1 , from the infection to the first consultation is unknown and must be estimated. It is often difficult to determine t_1 , because a patient may not recollect a tick bite, and because symptoms may not appear until late in the progress of the disease.

As the conditional probabilities for the consultation modules vary depending on t_1 , different values of t_1 will result in Bayesian networks with different quantitative specification. Hence, estimation of t_1 is a crucial part of the construction of the patient specific model. We tackle this problem by hypothesising a number of different models, and we identify those that best matches the evidence. Each model includes all consultations for the patient. Thus, we have a fixed structure with different instantiations of

the conditional probabilities for different values of t_1 .

The probability of a model given the evidence can be calculated by using Bayes' rule:

$$P(M | e) = \frac{P(e | M) \cdot P(M)}{P(e)} \quad (1)$$

$P(e | M)$ can be extracted from the hypothesised Bayesian network as the normalization constant. The probability of the evidence, $P(e)$, is not known, but is a model independent constant. The prior probability of the model, $P(M)$, is assumed to be equal for all models, and is therefore inversely proportional to the number of models. Thus, $P(M | e)$ can be expressed by

$$P(M | e) \propto P(e | M) \quad (2)$$

The probability of Lyme disease, $P(Ld)$ can be read from each model and will, of course, vary depending on the model M . By using the fundamental rule, the joint probability of $P(Ld, M | e)$ can be obtained as

$$P(Ld, M | e) = P(Ld | M, e) \cdot P(M | e) \quad (3)$$

where $P(M | e)$ can be substituted by using equation 2:

$$P(Ld, M | e) \propto P(Ld | M, e) \cdot P(e | M) \quad (4)$$

The probability of Lyme disease given the evidence, can now be found by marginalizing over M :

$$P(Ld | e) = \sum_{T_n=T_{start}}^{T_{end}} P(Ld | M_{T_n}, e) \cdot P(e | M_{T_n}) \quad (5)$$

T_{start} and T_{end} represent an interval that surrounds the most probable value of the time since the bite, t_1 , as illustrated in Figure 8.

The interval is chosen due to computational considerations, and because the assumption that all models are equally probable is obviously not valid. Alternatively, we could simply choose the model with highest probability.




Figure 8: The figure illustrates the probability of the different models as a function of days since infection.

The estimation of the most probable model has been implemented by evaluating all models for t_1 in the interval 0-800 days. The upper limit of 800 days is based on the assumption that the general development pattern of Lyme disease at this time is stabilized and the symptoms are chronic. The size of the interval $T_{start} - T_{end}$ has, somewhat arbitrarily, been set to 25 days.

5 Preliminary evaluation

In a preliminary test of the proposed approach the results were compared to the outcome of the model by Dessau and Andersen. The evaluation was focused on the structural problems in the models. The cases that were used for the evaluation are based on a survey performed by Ram Dessau where 2414 questionnaires have been collected from the Danish hospitals in Aalborg and Herlev. The overall picture of the comparison is that the proposed model in general estimates a higher probability of Lyme disease than the model by Dessau and Andersen.

In order to evaluate the effect of the incorporation of the clinical history in the time sliced model, a number of typical courses of Lyme disease have been evaluated as single consultations without utilizing the clinical history and as consultations utilizing the previous history. At the same time, each of the consultations have been evaluated in the model by Dessau and Andersen in order to evaluate how it handles the later stage symptoms, when the clinical history is not incorporated.

This informal study indicates that the time sliced model behaves as intended when the clinical history is incorporated. The estimated prob-

ability of Lyme disease is gradually increased for each added consultation, as the general development pattern of the disease is confirmed in typical courses.

In the estimates from the model by Dessau and Andersen the probability of Lyme disease decreases as the later stage symptoms are observed. This reasoning does not seem appropriate when it is known from earlier consultations that the clinical history points toward a borrelial infection. The model by Dessau and Andersen is not designed to incorporate the clinical history, but from the results it can be seen that this parameter is important in order to provide reasonable estimates of the probability of Lyme disease.

6 Conclusion

Temporal reasoning is an integral part of medical expertise. Many decisions are based on prognoses or diagnoses where symptoms evolve over time and the effects of treatments typically become apparent only with some delay. A common scenario in the general practice is that a patient attends the clinic at different times on a non-regular basis.

Lyme disease is an infection characterised by a number of symptoms whose manifestations evolve over time. In order to correctly classify the disease it is important to include the clinical history of the patient. We have proposed a method to include consultations scattered over time at non-equidistant points. A description of the evolution of symptoms over time forms the basis of a dynamically tailored model that describes a specific patient. The time of infection of the disease is estimated by a model search that identifies the most probable model for the patient given the pattern of symptoms over time.

Based on a preliminary evaluation it can be concluded that the method proposed for handling nonequivalent time intervals in order to incorporate the clinical history works satisfactory. In cases where the clinical history confirmed the general development pattern of Lyme disease the estimated probability the disease was

gradually increased from consultation to consultation, whereas it was reduced when the history did not confirm the pattern.

We conclude that the proposed method seems viable for temporal domains, where changing conditional probabilities can be modeled by continuous functions.

References

- [Arroyo-Figueroa and Sucar1999] Gustava Arroyo-Figueroa and Luis Enrique Sucar. 1999. A Temporal Bayesian Network for Diagnosis and Prediction. *In Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence*, pages 13–20. ISBN: 1-55860-614-9.
- [Dessau and Andersen2001] Ram Dessau and Maj-Britt Nordahl Andersen. 2001. Beslutningsstøtte med Bayesiansk netværk - En model til tolkning af laboratoriesvar. Technical report, Aalborg Universitet. In danish.
- [Fisker et al.2002] Brian Fisker, Kristian Kirk, Jimmy Klitgaard, and Lars Reng. 2002. Decision Support System for Lyme Disease. Technical report, Aalborg Universitet.
- [Gray et al.2002] J.S. Gray, O. Kahl, R.S. Lane, and G. Stanek. 2002. *Lyme Borreliosis: Biology, Epidemiology and Control*. Cabi Publishing. ISBN: 0-85199-632-9.
- [Jensen2001] Finn V. Jensen. 2001. *Bayesian Networks and Decision Graphs*. Springer Verlag. ISBN: 0-387-95259-4.
- [Nodelman et al.2002] Uri Nodelman, Christian R. Shelton, and Daphne Koller. 2002. Continuous Time Bayesian Networks. *In Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, pages 378–387. ISBN: 1-55860-897-4.
- [Smith et al.2003] Martin Smith, Jeremy Gray, Marta Granström, Crawford Revie, and George Gettinby. 2003. European Concerted Action on Lyme Borreliosis. <http://vie.dis.strath.ac.uk/vie/LymeEU/>. Read 04-12-2003.