How clear is transparent?
Reporting expert reasoning in legal cases

Marjan J. Sjerps[a,b], Charles E.H. Berger[a,c]

[a] *Netherlands Forensic Institute, The Hague*
[b] *University of Amsterdam*
[c] *Leiden University*

Experts providing evidence in legal cases are universally recommended to be transparent, particularly in their reasoning, so that legal practitioners can critically check whether the conclusions are adequately supported by the results. However, when exploring the practical meaning of this recommendation it becomes clear that people have different things in mind. The UK appeal court case *R* v *T* painfully exposes the different views. In this paper we argue that there can be a trade-off between clarity and transparency, and that in some cases it is impossible for the legal practitioner to be able to follow the expert's reasoning in full detail because of the level of complexity. All that can be expected in these cases is that the legal practitioner is able to understand the reasoning up to a certain level. We propose that experts should only report the main arguments, but must make this clear and provide further details on request. Reporting guidelines should address the reasoning in more detail. Legal practitioners and scientists should not be telling each other what to do in the setting of a legal case, but in other settings more discussion will be beneficial to both. We see the likelihood ratio framework and Bayesian networks as tools to promote transparency and logic. Finally, we argue that transparency requires making clear whether a conclusion is a consensus and reporting diverging opinions on request.

1. Introduction

Transparency is generally considered an essential requirement for expert evidence used in legal cases, in many kinds of legal systems and in many countries. The concept may refer to various aspects of expert evidence. For example, Mnookin et al. (2011) use it in relation to access to data from forensic science institutes for research purposes and to forensic examiners as test objects, and in acknowledging errors. In this paper, we restrict ourselves to a single different aspect: the expert's reasoning process.

In a recent UK court case, *R* v *T* (*R* v *T* 2010), the transparency of the expert's reasoning was one of the many issues that was criticized by the court. The problem was that the expert, a shoe print expert, had presented his evidence without mentioning that he had performed a likelihood ratio calculation (e.g. Robertson and Vignaux 1995, Taroni and Aitken 2004) based on imperfect data. In the development of the appeal, a reviewer of the expert's report noted that likelihood ratio calculations were performed but not reported. The expert's explanation for the omission was that it was no standard practice for the details relating to statistics and likelihood ratios to be included in a report. He made clear that the data were not available to an exact and precise level and it was only used to confirm an opinion substantially based on his experience and so that it could be expressed in a standardised form (*R* v *T*, para 38).

Indeed, the standards set out for reporting expert opinion by the AFSP (2009) state:

> *"Transparency - The expert will be able to demonstrate how he came to his conclusion. He will set out in the statement or report the basis of his opinion viz.:*
>
> - *Propositions addressed.*
> - *Test or examination results.*
> - *The background information he has used in arriving at his conclusion.*
> - *He will be able, if required, to provide the data he has used and its provenance."*

Note that this guideline recommends reporting the input of the reasoning process (propositions, results, background information, and data), but does not say anything about reporting the reasoning itself: e.g. why these propositions are chosen, or how the input leads to

the conclusion. The first sentence merely states that the expert should be able to explain the reasoning, but doesn't say if or how it should be reported. This is like giving someone the ingredients of the cake, but not the recipe to bake it. The same vagueness concerning reporting expert reasoning is found in other guidelines, e.g. the recent ISFG recommendations on forensic animal DNA (Linacre 2011) only states "The method of the evaluation of the weight of the evidence should be stated clearly…". In this paper we will propose a clarification of this point for the AFSP guideline.

The court of appeal in *R* v *T* regarded the omission of the likelihood ratio calculations as a serious lack of disclosure:

> *"The process by which the evidence was adduced lacked transparency. This is no personal criticism of Mr [...], as he was simply following practice. However, it is simply wrong in principle for an expert to fail to set out the way in which he has reached his conclusion in his report*." (R v T, para 108).

This view merits further discussion. It is fairly common for experts to report the basis of their opinion, and the arguments underlying the conclusion, but only up to a level that they think the readers of the report are able to understand. The omission of technicalities and minor details in the reasoning is broadly accepted and even promoted. Hence likelihood ratio calculations and other kinds of analysis might be performed to study the case, without being reported. To report it or not is a choice that is made for every individual case, and we will return to the choice made in this specific case later. An increasingly popular way to facilitate expert reasoning is the use of probabilistic graphical models called Bayesian networks (e.g.,Taroni et al. 2006). However, many experts are reluctant to present the Bayesian network at court.. For instance, Evett et al (2002) write concerning the use and presentation of Bayesian networks:

> *"We anticipate that Bayesian networks will play an increasingly important role in forensic science. This does not necessarily mean the vision of the scientist presenting networks at court — but their power in enabling the scientist to understand the fundamental issues in a case and to discuss*

them with colleagues and advocates is something that has not previously been seen within forensic science."

We note that consumers and producers of expert evidence agree about the necessity of transparency in expert reasoning, while having differing views on how this should be put in practice. We think that the cause of this difference lies in the rather vague nature of the concepts "transparency" and "expert reasoning". The purpose of this paper is to explore these concepts and to promote discussion about it. We will consider likelihood ratio calculations and the use of Bayesian networks in particular, as these are becoming increasingly important for expert reasoning. We do not claim to present original solutions or ideas. Many of our suggestions may well have been already implemented at various institutes. Our goal is to promote the exchange of ideas about the topic, so that guidelines addressing it can be more specific, and confusion in legal cases between experts and legal practitioners such as in R v T can be prevented.

We will first define what we mean with "expert reasoning", and "transparency" of it. Subsequently, we will consider the question whether we want the expert to be fully transparent. We will argue that full transparency cannot always result in full understanding, and that there can be a trade-off with clarity. Given these restrictions, we will arrive at what should be reported and what not. We will also consider the interaction between legal participants and experts. Finally, we will consider how current practice can be improved.

2. Definition of concepts

Expert reasoning

With the term "expert reasoning" we mean the interpretation or logical argumentation of the expert, showing how the conclusion follows from the results of an examination. This interpretation is the output of a thought process. The input of this process consists of the expert's knowledge and experience concerning similar earlier cases, literature, experiments, and of course the results in a current case. The case may also be discussed with colleagues, and their knowledge may be used as input. The process itself consists of deriving the questions to be answered, selecting from the input the information needed to answer the questions, structuring this information into a coherent argumentation for answering the questions, and formulating the conclusion.

The thought process may be made explicit in a likelihood ratio analysis, either in numerical or verbal form (see e.g. Cook et al 1998, Evett 1998) or in a graphical way using a Bayesian network (e.g. Aitken and Gammerman 1989, Taroni et al. 2006). The thought process requires the estimation of all sorts of uncertain parameters, e.g. frequencies, transfer probabilities, sampling and measurement uncertainties, sensitivities, etc. The accuracy of these estimates depends on the quality of the input information and may vary from a rough guess to a precise number. By combining all these estimates, either explicitly in a likelihood ratio calculation or implicitly, the expert reaches a conclusion. Subsequently, there are several ways of verifying, for instance by a crude calculation to check the order of magnitude of the likelihood ratio, performing real or thought experiments, running computer simulations, trying to falsify the conclusion with e.g. counter examples, and comparing with the conclusion of another expert who analysed the same case.

In *R* v *T* for example, the expert reached a conclusion implicitly, based on his experience. He also performed a likelihood ratio analysis that forced him to make this reasoning more explicit, even if the numbers were subjective, rough estimates or based on a database that is not fully representative. According to the expert, he used this analysis as a confirmation of his implicit analysis.

Transparency of expert reasoning

What do we mean when we require the expert to be transparent about his reasoning? As described above, expert reasoning is a sometimes complicated process involving a number of aspects such as deriving questions, selecting and structuring information, and formulating and verifying a conclusion. Transparency of reasoning in our view concerns the disclosure of those aspects. When writing an expert report, the following choices are made:

- Which aspects to report. The expert can choose to report all aspects, or make a selection and report only those aspects that are deemed relevant. It is also still fairly common for experts to report no aspects at all: the report simply jumps from results to conclusion.
- For each reported aspect A, the level of detail to report it. This level may vary from full detail, where every (technical) detail of A is reported, to no detail, where A is only mentioned.

The expert who reports every aspect of his reasoning in full detail is completely transparent about his reasoning. But do we really want the expert to do this?

3. Do we want experts to be fully transparent?

To answer this question we have to look at the purpose of the report. In order for legal practitioners to make effective use of expert reports in the legal process they must be able to understand what the expert concluded and how. They must also be able to be critical about it in the sense that they (and/or a counter-expert) should be able to judge whether the expert addressed the appropriate questions, whether the technical methods he used were sufficiently reliable, and whether the results support the conclusion. It is the latter that requires the expert to be transparent about the reasoning process. Obviously, the legal practitioner could do a better job if he were able to follow every detail in this process. But is this feasible?

Can full transparency result in full understanding?

The reasoning process leading to a conclusion often requires more than just "common sense" and basic logic. The forensic results often require knowledge about the process that produced them. But apart from this, the interpretation process involves probabilistic reasoning, as shown over the past decades in a line of research called "evidence interpretation". Forensic experts are (or should be) trained in this kind of reasoning, which can range in complexity from simple assessments that are easy to follow to rather complicated formulas or graphical models called Bayesian networks that require a lot of explanation when presented.

For the simple probabilistic assessments, it is perfectly feasible for legal practitioners and jury members to follow the reasoning in detail. For example, if the DNA profile of a crime stain does not match the suspect's profile, the DNA is from someone else (unless an unlikely error was made). The real challenges are the complex lines of probabilistic reasoning. It is our opinion that these can involve expert knowledge in the same way that the technicalities of the method that produced the results involve expert knowledge. It is an illusion to think that the legal practitioner should be able to follow such a reasoning process in full detail in a limited amount of time. All that can be expected in these cases is that the legal practitioner is able to understand the reasoning process up to a certain level. Beyond that level, he has to rely on the expert, his own counter expert, or be prepared to spend a significant amount of time and effort. That level should be such that he can judge whether a counter expert should be hired, and if not, to be able to critically consider the expert's reasoning himself. As Aitken et al. (2010) put it, "In short, judges, lawyers

and expert witnesses should be responsible producers and discerning consumers of statistical information and probabilistic reasoning whenever they are introduced into criminal proceedings."

For example, when a detailed and complicated Bayesian network such as presented in Cowell et al (2011) is used to analyse a DNA mixture, we cannot expect a legal practitioner or jury member to fully understand the details of this model, as this would entail knowledge about e.g. Dirichlet distributions and hyperparameters. However, the legal practitioner could follow the reasoning at a more general level and, guided by the expert, just look at the graph and see what is the input, what is the output, realise that there is a hidden layer with probability tables, get a general idea of robustness, and know the main assumptions and choices that were made. A (counter) expert should know every detail of the model, including the underlying probability tables, sensitivity analysis etc. To draw an analogy, a pilot (the legal practitioner) needs to basically understand how the plane works in order to recognise problems and fly it safely, and the plane technicians (the forensic experts) need to know the inner workings of the plane.

Our view that understanding the experts reasoning process is only feasible up to a certain level seems to be shared by others. Fenton and Neil (to appear) write: "While many members of the legal profession are aware of the fallacies, they struggle to understand and avoid them. This seems to be largely because they cannot follow Bayes' Theorem in its formulaic representation. Instead of continuing the painful struggle to get non-mathematicians to understand mathematics we must recognise that there is an alternative approach that seems to work better." They propose to present Bayesian networks in reports and in court as a way to explain the reasoning without using mathematics. Their approach allows legal participants to understand the logic of the reasoning, but not the mathematical proof of this logic.

A trade-off between clarity and transparency

We call the expert's reasoning clear when a lay person quickly understands the line of reasoning and the crucial parts in it. We call the reasoning transparent when all its aspects are addressed in detail. Can we achieve both? Unfortunately, as many experts have experienced at some point in their career, discussing a scientific analysis in full technical detail in a report or in court conflicts with clarity: most experts will be considered incomprehensible by legal practitioners and jury members. Furthermore, there is a limited amount of time that can be spent on the scientific evidence, so full transparency is not an option for practical reasons too. Consequently, one of the main tasks of an expert lies in the filtering of information: what is important in the case, and to whom?

Clarity is optimised when only the most important aspects of the experts reasoning are summarized at a lay person's level. Unfortunately, this might not be an option either because it can reduce transparency: legal processes can be unpredictable, and details that seemed unimportant at first can become crucial at a later stage. Also, the defence should be given the opportunity to question the filtering, and this requires disclosure of all information including those parts that the expert considers irrelevant.

A lazy and surprisingly popular option is to mention no aspect at all, and jump straight from results to conclusion. Surely this leads to the shortest, but also the worst report as this does not allow legal practitioners to check whether the results indeed support the conclusion.

Thus, legal participants can only understand the expert's reasoning up to a certain level, even when the expert is fully transparent. Furthermore, there is a trade-off between transparency and clarity. A good expert therefore filters information for the sake of clarity, while still providing all of it for the sake of transparency. In the following we will address the question how this can be achieved.

4. How to report expert reasoning?

Dealing with minor aspects

As described above the reasoning process of the expert has many aspects, and a good expert reports only the important ones. How to deal with the other aspects? Because of the trade-off between clarity and transparency, we think that the expert does a bad job if he considers these parts in the report. Most of the time it will just cloud the issues that really matter. We believe the solution lies in reporting that some parts of the reasoning are omitted, but are available on request. Quality assurance systems usually require that the underlying reasoning is documented completely in the case file, for further reference. The expert could state, for example, that for the sake of clarity the report only presents the main arguments underlying his opinion, and that further details are documented in the laboratory case file and are available on request. Another option is to present the "filtered" information in an appendix, or to structure the report or presentation otherwise to make a clear distinction between the relevant information and the information that the expert at that stage considers of minor importance for the legal process. In this way, the expert's reasoning is clear, and complete in the sense that all aspects are disclosed, just not all aspects are explained.

Which aspects should be reported?

We concluded that only aspects that are of interest should be reported. But when are they of interest? This is a difficult question in general. We think that for answering it, it is useful to distinguish between arguments and the tools used to find or explore them. To clarify this, let us first consider a simple example. Suppose that an expert discusses a case with a colleague over lunch, and that this colleague suggested considering the "missing" evidence: evidence that was not observed but which was expected under the prosecution hypothesis. In this example, the argument is the "missing" evidence, and the tool to find it was the colleague during lunch. For the legal process it is important that the expert considers the "missing" evidence in his report, but mention of the colleague is superfluous and it is safe to only document the discussion with him in the case file.

Now let us consider a more complicated example. In a fibre case the question of interest is how a set of fibres was transferred, by direct contact or otherwise. The expert will base his opinion on many arguments concerning transfer, persistence and recovery of the fibres, e.g. the probability of transfer from the suspect's wool sweater to the victim's clothes in case of direct contact is large because this type of wool easily sheds many fibres. Suppose that this probability is slightly less if the sweater is wet. The expert may construct a Bayesian network model, which requires thinking about which hypotheses may be addressed, which variables are relevant, their dependencies, and which assumptions are necessary. Based on the graphical model, the software automatically shows which probabilities need to be estimated. The construction of the network thus provides insight in the problem and serves as a visual aid to structure thought, making sure he does not overlook considering some variable, dependency, probability or assumption. The network also facilitates discussion between experts as a visual representation of the reasoning that quickly shows what was taken into account. The expert may thus use the network as a tool to produce a complete and coherent reasoning in writing. For the lawyer, however, it is not of interest to know what tool was used to facilitate the reasoning, a Bayesian network, a discussion over lunch, or anything else. Only the reasoning itself is of interest. It suffices to document the Bayesian network in the case file (and report, as suggested above, that only the main arguments are reported and further details are available from the case file).

The network may subsequently be used as a probability calculator to produce e.g. a numerical likelihood ratio, or a sensitivity analysis. However, the output is only as reliable as the input of the network (probabilities of transfer, recovery and persistence of wool fibres in certain circumstances). When the input data are sufficiently reliable, the expert may use the numerical

output of the network as part of the reasoning, or even as the conclusion itself. In DNA relatedness analysis for instance, the reported "paternity index" is in fact a likelihood ratio. For the lawyer, it is of interest to know how such numbers were produced, including important details such as assumptions made (Gjertson et al. 2007).

Generalizing the above, we conclude that experts use certain tools to arrive at a set of main arguments supporting their conclusion. In our opinion, it is this set that is important and that should be reported, rather than the tools used to arrive at it. Consequently, when a method is sometimes used as a tool to facilitate reasoning, and sometimes as an argument, it should only be reported in case it is used as an argument. We therefore believe that demanding that certain methods should always be, or not be, reported (like the court in *R* v *T* does for likelihood ratio methods) is not in the interest of the legal process. The selection criterion for reporting a certain aspect of the expert's thought process should be the added value to the argumentation. Since assessing this added value typically requires scientific knowledge, (and also for practical reasons) it is the scientist who must decide on this. It may always happen that a minor argument that was not reported becomes important in the case. The only solution to prevent this from happening is demanding full transparency. We have explained above why this is not a good solution.

For instance for the fibre case, it is up to the scientist to decide that the network is just a tool for structured thinking that is documented in the case file but isn't worth reporting. It is also up to him to select from all the arguments the main ones that need reporting (e.g., the large transfer probability of the wool fibres) and the minor ones that are documented in the case file only (e.g., that this probability is a bit smaller if the sweater was wet). It may turn out later in the case that the sweater was soaked. If the fibre evidence is crucial, the lawyer can ask for the details in the case file, if it is made clear in the report that such details are available.

Should the expert have reported the likelihood ratio calculation in *R* v *T*? The Court considered the omission a serious lack of disclosure. They suspected that the reason for the omission was that "outside the field of DNA (and possibly other areas where there is a firm statistical base) this court has made it clear that Bayes theorem and likelihood ratios should not be used." (*R* v *T*, para 90), and "The justification advanced in the evidence in this appeal for not including in the reports the use of the formula and statistics was that it might confuse the jury. No doubt this was a reaction to the perceived consequences of the views of this court expressed in Adams and subsequent cases." (*R* v *T*, para 99). The expert explained why he considered the likelihood ratio calculation of minor importance: "He made clear that the data were not available to an exact and precise level and it was only used to confirm an opinion substantially based on his experience and so that it could be expressed in a standardised form." (*R* v *T*, para 38). According

to the expert, his (experience-based) arguments concerned pattern, size, wear and damage of the shoes, and the likelihood ratio calculation was merely used as a tool to explore and confirm these arguments. In our view, if the expert truly considered the added value to his experience based reasoning as negligible, he made the correct decision not to report it. This raises the question to what extent the numerical value of the likelihood ratio determined the level of the verbal conclusion ("moderate support"). In any case, he could have stated that his opinion was based on a number of sources of information, that he reported the main arguments and that further details were available on request. Perhaps a simple statement like this could have prevented the impression that there was (an intentional) failure of disclosure.

What is the appropriate level of detail?

In the above we saw that the degree of transparency about a certain aspect can vary from no to full detail. Obviously, counter experts should be provided with full details. For legal practitioners, the appropriate level of detail increases with the importance of the argument. For crucial arguments experts should aim for the highest level of detail that is still understandable for legal practitioners. Thus, if a likelihood ratio formula, a frequency calculation based on a database, or a Bayesian network is the foundation of the expert's opinion, the expert should try to convey all information needed to explain the reasoning but avoid mathematical or technical details that distract from the argument rather than explain it. Fenton and Neil (to appear) are also in favour of avoiding mathematical details: "Crucially, there should be no more need to explain the Bayesian calculations in a complex argument than there should be any need to explain the thousands of circuit level calculations used by a calculator to compute a long division. Lay people do not need to understand how the calculator works in order to accept the results of the calculations as being correct to a sufficient level of accuracy. The same must eventually apply to the results of calculations from a BN tool. […] Because of this, we believe that in 50 years time professionals of all types involved in the legal system will look back in total disbelief that they could have ignored these available techniques of reasoning about evidence for so long." Time will tell if this view is realistic. We agree that explanation would more likely be concerned with the structure of the network and its probability tables than with computational details.

As an example, when a detailed and complicated Bayesian network is used to analyse e.g. a DNA mixture, the report could just present the important assumptions and the argumentation derived from the network analysis in a verbal way. The expert could then mention that his opinion is based on a probabilistic model, with details available on request. When more details are

requested, the network can be presented but not necessarily in full mathematical detail. Fenton and Neil's (to appear) approach can be used to do this. It should always be clear, however, that full details are provided when asked for or when counter experts are involved.

5. Interaction between legal practitioners and experts

Can the legal practitioner tell the scientist how to do science?

In R v T (para 95) the Court states: "In our judgment, an expert footwear mark examiner can therefore in appropriate cases use his experience to express a more definitive evaluative opinion where the conclusion is that the mark "could have been made" by the footwear. However no likelihood ratios or other mathematical formula should be used in reaching that judgement for the reasons we have given." We cannot interpret this statement other than legal practitioners telling the scientist how to do science, which we regard to be unacceptable. Legal practitioners would be shooting in their own foot if they would succeed in this interference with science. Not surprisingly, there was a protest against both the principle and content of this judgment, supported by many forensic scientists [e.g., position statement 2011, Berger et al. 2011). In our opinion, the expert should be transparent in his reasoning, allowing the legal practitioner to be critical about the expert's opinion and either use it in the legal process or not. However, when criticised in court it is up to scientists, and not legal practitioners, to decide whether the criticism is scientifically relevant and the reasoning in subsequent cases should be adjusted. Neufeld and Scheck (2010) also warn against the interference of legal practitioners with science: "Political and criminal justice ends -rather than research imperatives- have taken forensic science off course. We can bring about substantial improvement only if we trust scientists to take charge of forensic science, free of politics and independent of outside influence."

Can the scientist tell the legal practitioner how to practice law?

Interestingly, the reaction of the Court in *R* v *T* may have been triggered by a previous case in which a scientist showed how the jury could combine the scientific with the non-scientific evidence using elementary probability calculus. When discussing this well-known Adams case, Evett and Weir (1998, page 238) conclude: "The root problem is that the legal system exerts powerful forces against carrying through the most appropriate procedure effectively." Indeed, the Court in *R* v *T* writes (para 46): "It was submitted to the court that the approach adopted was a Bayesian analysis which this Court had robustly rejected for non DNA evidence in a number of

cases: *R v Dennis Adams* [1996] 2 Cr App R 467, *R v Adams* (No 2) [1998] 1 Cr App 377; *R v Doheny* [1997] 1 Cr App R 369." Gill et al. (2006) see a responsibility for scientists to educate legal practitioners: "Recommendation 2: Even if the legal system does not implicitly appear to support the use of the likelihood ratio, it is recommended that the scientist is trained in the methodology and routinely uses it in case notes, advising the court in the preferred method before reporting the evidence in line with the court requirements. The scientific community has a responsibility to support improvement of standards of scientific reasoning in the court-room."

We think that scientists and legal practitioners should not tell each other how to do their job in the setting of a legal case. Professional discussion and interaction between the two groups in other settings however, is to be promoted. This holds for academics but also for practitioners. In the UK and in the Netherlands, for example, such interactions are currently emerging. The International Conference on Forensic Inference and Statistics is an initiative dating back to 1990 that aims to bring together forensic scientists / statisticians and legal practitioners. Fenton and Neil (to appear) conclude: "There is clearly a general problem for the statistics community about how to get their voices heard within the legal community. […] If we can similarly break down the barriers between the mathematics and law communities, then there could be a transformative impact in the way legal evidence is analysed and presented."

6. Current practice and room for improvement

Reporting expert reasoning

The way that expert reasoning is reported varies widely over countries, areas of expertise, and experts. Many reports simply jump from the "results" section to the "conclusion", without explaining how the results lead to this conclusion. Other reports include an "interpretation of results" section, which may be very lengthy and detailed but which may also be very short. Information sheets providing information about the method and the basis of the conclusion in general can accompany the report. For those reports that jump from results to conclusion, or that address the interpretation of the results only superficially, there is certainly room for improvement.

It is important not to confuse an agreed procedure relating results to conclusion with the reasoning underlying the conclusion. Often the procedure is clear, but the reasoning itself is not. This is the source of a common misunderstanding, where the expert may think he is perfectly transparent in how he arrived at his conclusion, whereas others may criticise his method as a black box. For instance, when we consider the field of fingerprint evidence, in many countries the

procedure by which a conclusion is reached is crystal clear: a number of corresponding minutiae is identified, and the conclusion follows from this number and a threshold. However, it is the subjective nature of the identification of the minutiae, and the underlying logic of the threshold and the conclusion that is the problem (e.g. Cole 2008).

Improving (transparency of) expert reasoning

There is a general sense that there is room for improvement in the reasoning itself: many scientists feel a fundamental change is necessary in forensic science, and the direction is towards the use of probabilistic and statistical methods (e.g. Lindley 1977, Evett 1998, Saks and Koehler 2005, Curran 2009, Saks 2010). The likelihood ratio framework for interpretation of evidence is increasingly used, and this will ultimately increase the transparency of expert reasoning. Concerning this claim in footwear mark evidence and the *R* v *T* case, Skerret et al. (2011) write: "...we have shown, using casework examples, how the model can act as a guide to increase the transparency of the evaluation and reporting of the weight of shoe mark evidence." Also, Bayesian networks are promoted as tools for creating transparency: "Detailed statistical modelling is practically and economically infeasible in such "messy" situations, but nevertheless judgements have to be made. In practice these decisions can be a black art, involving opaque assumptions and unchecked subjectivity, but in our experience Bayesian methods can help bring some rigor and structure. More importantly, they also encourage transparency and allow uncertainties and assumptions to be modelled explicitly." (Fenton and Neil 2006).

Indeed, in our view an expert who argues: "the shoe mark was probably left by the shoe, because based on my long-term experience I can tell that there are enough corresponding acquired features to justify this conclusion" is absolutely non-transparent: the argumentation is a logical black box. An expert who argues that the corresponding acquired features provide strong evidence to support the view that the shoe made the mark (versus that some other shoe with the same size and sole pattern made the mark), and presents a verbal argumentation along the lines of a likelihood ratio analysis (that is documented in the case file), is much more transparent: the logical argumentation can be explained in full detail if necessary.

Does transparency require reporting diverging opinions?

In many forensic science laboratories, each case is interpreted by two or more caseworkers independently. When opinions diverge, the currently predominant practice is to discuss the case with each other, formulate a consensus conclusion and report only this, without mentioning that originally opinions diverged. Sometimes consensus is reached easily because one of the experts discovers he had overlooked something. It may also happen that one expert convinces the other (hopefully not by using seniority as an argument). However, if courts and scientists truly value transparency then it is better to report all opinions separately, or report them as an appendix of the report. Mnookin et al. (2011) seem to support reporting diverging opinions: "Perhaps more robust reporting requirements, in which an analyst routinely discloses any interpretive disagreement within her laboratory report, would be a simpler means to achieve a similar goal." In our opinion, in most cases consensus opinions are valuable as a joint conclusion shared by experts but it is important to make clear that the opinion is a consensus, and that the original opinions are documented in the case file (available on request). The use of likelihood ratio analyses or Bayesian networks facilitates understanding at which points opinions diverge. Concerning Bayesian networks, Aitken and Gammerman wrote back in 1989: "There are many conceptual problems to be solved here but the procedure enables different opinions to be modelled quickly and coherently. Problems which may arise over assessment of these probabilities will be highlighted by the procedure."

Improving guidelines

Given the importance of transparency of expert reasoning, we propose to include this topic in reporting guidelines. In particular, the AFSP guideline (AFSP 2009) states:

> *"Transparency - The expert will be able to demonstrate how he came to his conclusion. He will set out in the statement or report the basis of his opinion viz.:*
>
> - *Propositions addressed.*
> - *Test or examination results.*
> - *The background information he has used in arriving at his conclusion.*

- *He will be able, if required, to provide the data he has used and its provenance."*

We propose to change the first sentence: "The expert will report how he came to his conclusion. He will report only those arguments that in his view are important. He will document his reasoning in full detail in the case file."

We also propose to add four bullets:

- Main assumptions
- The interpretation and evaluation of the results in the light of the propositions and the background information, thus explaining how the results lead to his conclusion (the likelihood ratio).
- He will make clear that only main arguments are reported and that full details are available on request from the case file.
- He will make clear if the opinion is a consensus opinion, and that the original opinions are available on request from the case file.

7. Conclusions and recommendations

We summarize our conclusions and recommendations below:

- We have defined expert reasoning as the argumentation showing how the conclusion follows from the results. It is the output of a thought process that involves many aspects (concerning deriving questions, selecting and structuring information, formulating and checking a conclusion).
- Transparency of reasoning concerns two decisions: which aspects to report and at what level of detail.
- Transparency is required because the legal practitioner needs to understand the expert's reasoning to check whether the results support the conclusion. However in complex cases, the legal practitioner can only understand the reasoning up to a certain level.
- There can be a trade-off between clarity and transparency.
- The scientist should only report those aspects of his reasoning that have significant added value to the argumentation. Other aspects should be documented in the case file but omitted from the report for sake of clarity.

- The expert should aim for a high level of detail only for crucial parts of the reasoning. Less important arguments can be presented at a more global level. The report should make clear that only the main arguments are presented and that further details are available on request.
- Legal practitioners and scientists should not tell each other how to do their job in the setting of legal cases, but discussion between the two groups in other settings is to be promoted.
- There is currently room for improvement in the expert's reasoning itself, and sometimes in reporting it. The likelihood ratio framework and Bayesian networks promote transparency and logic.
- Transparency requires making clear whether a conclusion is a consensus and reporting diverging opinions on request.
- We recommend to explicitly address transparency of expert reasoning in reporting guidelines.

References

AFSP (Association of Forensic Science Providers) (2009) Standards for the formulation of evaluative forensic science expert opinion, Science & Justice, 49,161–164.

Aitken C.G.G., Roberts P., Jackson G. (2010) Fundamentals of Probability and Statistical Evidence in Criminal Proceedings Guidance for Judges, Lawyers, Forensic Scientists and Expert Witnesses. London: The Royal Statistical Society.

Aitken C.G.G., Gammerman A.J. (1989) Probabilistic reasoning in evidential assessment, Journal of the Forensic Science Society, 29, 303–316.

Aitken C.G.G., Taroni F. (2004) Statistics and the evaluation of evidence for forensic scientists, 2nd ed., Chichester: Wiley.

Berger C.E.H., Buckleton J.S., Champod C., Evett I.W., Jackson G. (2011) Evidence evaluation: a response to the Appeal Court judgment in R v T, Science & Justice, 51, 43–49.

Cole S.A. (2008) Comment on 'Scientific validation of fingerprint evidence under Daubert', Law, Probability and Risk 7, 119–126.

Cook R., Evett I.W., Jackson G., Jones P.J. and Lambert J.A. (1998) A model for case assessment and interpretation, Science & Justice, 38,151–156.

Cowell R.G., Lauritzen S.L., Mortera J. (2011) Probabilistic expert systems for handling artefacts in complex DNA mixtures, Forensic Science International: Genetics, 5, 202–209.

Curran J.M. (2009) Statistics in forensic science, Wiley Interdisciplinary Reviews: Computational Statistics, 1, 141–156.

Evett I.W., Weir B.S. (1998) Interpreting DNA evidence: statistical genetics for forensic scientists, Sinauer Associates, Sunderland, Massachusetts, USA.

Evett I.W. (1998) Towards a uniform framework for reporting opinions in forensic science casework, Science & Justice, 38, 198–202.

Evett I.W., Gill P.D., Jackson G., Whitaker J., Champod C. (2002) Interpreting small quantities of DNA: the hierarchy of propositions and the use of Bayesian networks, J Forensic Sci., 47, 520–30.

Fenton N.E., Neil M., Avoiding Legal Fallacies in Practice Using Bayesian Networks', http://www.eecs.qmul.ac.uk/~norman/papers/fenton_neil_prob_fallacies_June2011web.pdf, website visited September 1st, 2011, according to website to appear in the Australian Journal of Legal Philosophy.

Gill P., Brenner C.H., Buckleton J.S., Carracedo A., Krawczak M., Mayr W.R., Morling N., Prinz M., Schneider P.M., Weir B.S. (2006) DNA commission of the International Society of Forensic Genetics: Recommendations on the interpretation of mixtures, Forensic Science International, 160, 90–101.

Gjertson D.W., Brenner C.H., Baur M.P., Carracedo A., Guidet F., Luque J.A., Lessig R., Mayr W.R., Pascali V.L., Prinz M., Schneider P.M., Morling N. (2007) ISFG: Recommendations on biostatistics in paternity testing, Forensic Science International: Genetics, 1, 223–231.

Honess T.M., Levi M., Charunan E.A. (1998) Juror Competence in Processing Complex Information: Implications from a Simulation of the Maxwell Trial, Criminal Law Review, Nov. 1998, 763–773.

Linacre A., Gusmaõ L., Hecht W., Hellmann A.P., Mayr W.R., Parson W., Prinz M., Schneider P.M., Morling N. (2011) ISFG: Recommendations regarding the use of non-human (animal) DNA in forensic genetic investigations, Forensic Science International: Genetics, 5, 501–505.

Lindley D.V. (1977) Probability and the Law, Journal of the Royal Statistical Society, Series D (The Statistician), 26, 203–220.

Mnookin J.L., Cole S.A., Dror I.E. , Fisher B.A.J., Houck M.M., Inman K., Kaye D.H., Koehler J.J., Langenburg G., Risinger D.M., Rudin N., Siegel J., Stoney D.A. (2011) The Need for a Research Culture in the Forensic Sciences, UCLA Law Review, 58, 725–779.

Neufeld P., Scheck B. (2010) Making forensic science more scientific, Nature, 464, 351–351.

Fenton N., Neil M. (2006) Comment: Expert Elicitation for Reliable System Design, Statistical Science, 21, 451–453.

Position statement (2011) Expressing Evaluative Opinions: a position statement, Science & Justice, 51, 1–2.

R v T, Court of Appeal – Criminal Division, EWCA Crim 2439, 2010.

Robertson B., Vignaux G.A. (1995) Interpreting Evidence: Evaluating Forensic Science in the Courtroom, Wiley, Chichester UK.

Saks M.J. (2010) Forensic identification: From a faith-based “Science” to a scientific science, Forensic Science International, 201, 14–17.

Saks M.J., Koehler J.J. (2005) The Coming Paradigm Shift in Forensic Identification Science, Science, 309, 892–895.

Skerrett J., Neumann C., Mateos-Garcia I. (2011) A Bayesian approach for interpreting shoemark evidence in forensic casework: Accounting for wear features, Forensic Science International, 210, 26–30.

Taroni F., Aitken C., Garbolino P., Biederman A. (2006) Bayesian networks and probabilistic inference in forensic science, Wiley, Chichester UK.