

A Probabilistic Constraint Satisfaction Model of Information Distortion in Diagnostic Reasoning

York Hagmayer (york.hagmayer@kcl.ac.uk)

Olga Kostopoulou (olga.kostopoulou@kcl.ac.uk)

Department Primary Care and Public Health Sciences, King's College London
42 Weston Street, London, SE1 3QD

Abstract

Information distortion is a cognitive bias in sequential diagnostic reasoning. Assumptions about the diagnostic validity of later evidence are distorted in favor of the leading hypothesis. Therefore the bias contributes to a primacy effect. Current parallel constraint satisfaction models account for order effects and coherence shifts, but do not explain information distortion. As an alternative a new, probabilistic constraint satisfaction model is proposed, which considers uncertainty about diagnostic validity by defining probability distributions over coherence relations. Simulations based on the new model show that by shifting distributions in favor of the leading hypothesis an increase in coherence can be achieved. Thus the model is able to explain information distortion by assuming a need for coherence. It also accounts for a number of other recent findings on clinical diagnostic reasoning. Alternative models and necessary future research are discussed.

Keywords: Diagnostic reasoning; information distortion; parallel constraint satisfaction model.

Information Distortion in Diagnostic Reasoning

Diagnostic reasoning is an important cognitive activity in many areas. Based on available evidence decision makers infer hidden properties or diagnoses that account for the observations made. Diagnostic reasoning is maybe most important in the clinical domain. Making accurate diagnoses is essential for physicians. Unfortunately clinical diagnostic reasoning is affected by many biases, which may result in medical error (Croskerry, 2003; Kostopoulou et al., 2008). One of these biases is *information distortion*. When deriving a diagnosis clinicians have been shown to bias their interpretation of newly arriving evidence to support their preferred hypothesis (Kostopoulou, Russo, Keenan, Delaney & Douri, 2012). More precisely, clinicians alter their assumptions about the diagnostic validity of observed signs and symptoms (i.e., the likelihood of the diagnosis given the sign) so that they lend greater support to the favored diagnostic hypothesis. Similar findings on pre-decisional distortion of evidence have been reported for other professions like sales (Russo et al., 2006). Information distortion has been explained by a need for coherence (Russo, Medvec, & Meloy, 1996). By interpreting new evidence as supportive of the leading hypothesis decision makers increase the coherence among the favored diagnostic hypothesis and the evidence. Consistency theories in turn account for the need for coherence (cf. Simon et al., 2004).

Parallel constraint satisfaction models, especially Thagard's (1989) ECHO model, have been used to implement coherence-based accounts of diagnostic reasoning (e.g., Gloeckner, Betsch & Schindler, 2009). These models were extended to sequentially arriving evidence, which affords frequent updating of diagnostic hypotheses (Mehlhorn, Taatgen, Lebriere, & Krems, 2011; Wang, Johnson, & Zhang, 2006). Although these models can account for biased decision making, they cannot fully explain information distortion. Constraint satisfaction models in general assume that coherence relations among evidence and hypotheses, which represent assumptions about diagnostic validity, are stable. Research on information distortion, however, shows that decision makers are uncertain about these relations and may change respective beliefs during decision making (Kostopoulou et al., 2012; Russo et al., 1996; 2006). To account for these findings we will put forward a new, probabilistic constraint satisfaction model.

In the paper, we will first briefly describe a recent study on information distortion to exemplify methods and findings. Then we outline a standard constraint satisfaction model of sequential diagnostic reasoning and discuss its shortcomings. Next we propose a new, probabilistic constraint satisfaction model. Results from a simulation study will show that the model predicts information distortion and other findings from the literature. Finally, alternative models will be discussed and necessary future research will be pointed out.

Exemplary Empirical Findings

Kostopoulou and colleagues (2012) recently published a study on information distortion in the clinical domain. Physicians were confronted with case vignettes presenting diagnostic evidence and asked to evaluate two competing diagnostic hypotheses A and B. Evidence was presented sequentially in a particular order. The first set of cues strongly favored Hypothesis A over B, the next set of cues equally supported both hypotheses, while the third set strongly favored Hypothesis B over A. Overall the evidence was ambiguous. Participating clinicians were asked to make two judgments after each new item: (i) to rate how much this particular item favors either hypothesis (i.e., the item's differential diagnostic validity), and (ii) to rate the likelihood of the diagnoses given all information received so far. Both ratings were made on a scale ranging from one hypothesis to the other. In addition, a control group of

physicians rated each item individually. Information distortion was calculated by computing the difference between individual cue ratings and mean control ratings. From a normative perspective, no information distortion should be expected as the diagnostic validity of individual cues should be constant. Hence, any changes in assumptions about diagnostic validity, which create additional support for the favored hypothesis, constitute a bias.

Three findings are important for the purpose of this paper (see Kostopoulou et al., 2012, for complete results). First, there was a substantial variation between clinicians with respect to the assumed diagnostic validity of cues, which indicates that clinicians were uncertain about how much each piece of evidence supported the hypotheses. Second, participants exaggerated or reduced the diagnostic validity of individual items to support the initially preferred hypothesis. This was especially true for the neutral cues. Third, a majority (56%) kept the initially preferred hypothesis, while 38% switched to the hypothesis favored by the evidence coming in last. Only 6% correctly judged the hypotheses as equally likely. A good model should be able to account these findings.

Constraint Satisfaction Models of Diagnostic Reasoning

There are many cognitive models to describe sequential hypothesis testing, including Bayesian and logical accounts. We focus on parallel constraint satisfaction models here as they have been very successful in modeling sequential diagnostic reasoning. They also account directly for the frequently found primacy, recency and coherence effects (Mehlhorn et al., 2011, Wang et al., 2006). Thirdly, they are supported by consistency theories, which provide a psychological plausible explanation for why people strive for coherence (Simon et al., 2004).

Many constraint satisfaction models are based on ECHO, a connectionist model of the theory of explanatory coherence (Thagard, 1989). The theory assumes that the acceptance of a belief depends on its relations to other beliefs. By accepting and rejecting beliefs, the overall coherence of the belief set can be maximized. Roughly speaking, a set of beliefs is coherent, if (i) beliefs connected by a positive link (i.e., mutual support, consistency or entailment) are both accepted or rejected, and (ii) only one of the beliefs connected by a negative link is accepted (see next paragraph for formal details).

ECHO has been implemented as a connectionist network (see Figure 1). Hypotheses and items of evidence are represented by nodes, while coherence relations are represented as symmetrical links. Hypotheses and evidence are connected by links with positive weights if they are coherent with each other (e.g., if the evidence is a diagnostically valid indicator), by negative links if they are incoherent (e.g., if the evidence indicates the absence of the diagnosis), or they are not related if they are irrelevant for each other. Evidence nodes are assumed to have a special status as their acceptance not only depends on coherence with

other beliefs but on observations. Therefore they receive external activation from a special activation unit (not shown in Figure 1). Evidence nodes in turn activate potential diagnoses. Hypotheses coherent with the evidence get positive activation, while contradicted hypotheses are negatively activated. Different hypotheses are assumed to compete in explaining the observations. Therefore they are negatively related. Activations are passed through the network and added to each other until a stable state is reached. More precisely, the activation of each unit j is updated by combining its current activation $a_j(t)$ with the net effect (net_j) of all the units i connected to it according to the following formalism (Thagard, 1989; see also McClelland & Rumelhart, 1981):

$$a_j(t+1) = a_j(t) * (1 - \Theta) + (net_j * [max - a_j(t)], \text{ if } net_j > 0 \\ = a_j(t) * (1 - \Theta) + (net_j * [a_j(t) - min], \text{ if } net_j \leq 0 \\ \text{with } net_j = \sum_i rel_{ij} * a_i$$

The parameter Θ represents a decay and min and max the maximum and minimum activation (usually 1 and -1). Final activations represent acceptance. Hence, the hypothesis, which receives the highest positive activation in the end, is the preferred diagnosis.

The coherence of a belief set can be calculated by summing up the products of final activations and relations. This measure has been called harmony (Thagard, 1989).

$$\text{Harmony} = \sum_i \sum_j rel_{ij} * a_i * a_j$$

To account for sequentially arriving evidence, the external activation of evidence nodes is assumed to shift towards the new arriving evidence (Wang et al., 2006). In line with findings on the limited capacity of attention, the received activation is kept constant and is decayed exponentially across items. The activation received by an item of evidence is calculated according to the following equation:

$$a_{ev} = a_{ev} * \exp(-1 * \Lambda * [\text{Number of subsequent items seen}])$$

The parameter Λ represents how strongly the activation of an item is decayed due to later items. Little or no decay results in primacy effects, i.e., the first evidence biases decisions in favor of the initially preferred hypothesis. Strong decay leads to recency effects (Wang et al., 2006).

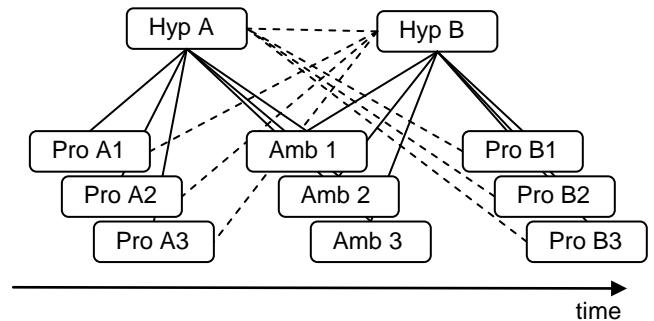


Figure 1. Parallel constraint satisfaction model of sequential diagnostic reasoning. Nodes represent hypotheses (HypA/B) and pieces of evidence (e.g., ProA). Solid lines represent coherent, dashed lines incoherent relations. Pieces of evidence arrive sequentially along the time line.

Figure 1 shows the structure of a constraint satisfaction network with two competing hypotheses (Hyp A, Hyp B) and nine pieces of evidence. The first three observations (ProA 1-3) support Hypothesis A (indicated by the solid lines) and contradict Hypothesis B (indicated by the dashed lines). The next observations (Ambig 1-3) support both hypotheses, while the final set favors Hypothesis B over A (ProB 1-3). This is the order of evidence clinicians received in the study by Kostopoulou et al. (2012).

This model predicts that Hypothesis A will be favored over Hypothesis B unless there is a very strong decay of the initial evidence (see simulations for respective evidence). But it cannot explain information distortion. As outlined above, information distortion means that assumptions about the diagnostic validity, i.e., the relations between evidence and hypotheses are distorted. The model presented here keeps these relations constant assuming that decision makers have stable beliefs about coherence relations. Hence the model cannot account for the findings by Kostopoulou and colleagues (2012) that participating clinicians distorted their assumptions about diagnostic validity for a particular case depending on their currently favored hypothesis.

A probabilistic constraint satisfaction model

Decision makers may be uncertain about the coherence relations among evidence and hypotheses. Consider the medical domain. Although a particular diagnostic cue may have a positive predictive value for Diagnosis A, there will be cases in which another diagnosis will prove to be true. Standard constraint satisfaction models of diagnostic reasoning do not allow us to represent this uncertainty. This uncertainty can be captured by conceptualizing the relations connecting evidence and hypotheses as beliefs and defining probability distributions over these beliefs. Probability distributions are used to represent the uncertainty in many cognitive models, e.g., Bayesian models (Chater & Oaksford, 2008), but they have not been used in constraint satisfaction models so far. Nevertheless, their application seems straightforward. There are three types of coherence relations: positive links, negative links, and no links (representing irrelevance). The probability distribution defines the likelihood that evidence and hypothesis are connected by a positive, a negative or no relation.

For example, to represent the assumption that a piece of evidence X almost always supports a Hypothesis A the probability of a positive link between X and A is set to a high value (i.e., $P(+_{AX}) \approx 1$) while the probabilities of a negative or no link are assumed to be very small (i.e., $P(-_{AX}) \approx 0$, $P(0_{AX}) \approx 0$). To derive predictions for a particular probability distribution, a set of constraint satisfaction networks is instantiated and run. Based on the resulting activations of the networks the likelihood that each hypothesis will receive the highest final activation is estimated. In addition, the mean resulting harmony is calculated to estimate the expected overall coherence.

Like standard parallel constraint satisfaction models the probabilistic models can account for primacy and recency effects by assuming differential decay of sequentially arriving information. Moreover, they may also account for information distortion. By shifting the probability distribution over coherence relations the overall coherence (i.e., harmony) may be increased. Thus a need for coherence may cause a change in beliefs about coherence relations resulting in information distortion. There is a limit however. To preserve the belief that a certain piece of evidence is coherent with a hypothesis in general, the probability distribution can only be shifted to a certain degree. To be more precise, the sign of the sum of weights of the relations multiplied with their respective probabilities has to remain the same. For example, if Hypothesis A and Evidence X are assumed to be coherent $\sum P(\text{relation}_{AX}) * \text{relation}_{AX} > 0$. Thus probabilistic constraint satisfaction models may predict information distortion without assuming an outright change in beliefs about the diagnostic validity of cues.

Simulations

To explore the predictions of probabilistic constraint satisfaction models, we implemented the model shown in Figure 1 with various probability distributions over coherence relations (see Table 1). The overall relation between each piece of evidence and the two hypotheses was kept the same across all distributions. The first three pieces of evidence were generally coherent with Hypothesis A and incoherent with B, the ambiguous evidence supported both hypotheses, and the final set contradicted A and favored B.

Table 1: Probability distributions over coherence relations of the model depicted in Figure 1.

Relation	HypA – Pro A1-A3			HypB – Pro A1-A3			HypA – Amb1-Amb3			HypB – Amb1-Amb3			HypA – Pro B1-B3			HypB – Pro B1-B3			
	Distrib.	P(+)	P(-)	P(o)	P(+)	P(-)	P(o)												
M1 fixed		1	-	-	-	1	-	1	-	-	1	-	-	-	1	-	1	-	-
M2		.9	-	.1	-	.9	.1	.9	.05	.05	.9	.05	.05	.05	.9	.05	.9	.05	.05
M3 ProA		.9	-	.1	-	.9	.1	.9	.05	.05	.5	.25	.25	.3	.4	.3	.4	.3	.3
M4 ProB		.9	-	.1	-	.9	.1	.5	.25	.25	.9	.05	.05	.05	.9	.05	.9	.05	.05
M5		.5	-	.5	-	.5	.5	.5	.25	.25	.5	.25	.25	-	.5	.5	.5	-	.5
M6 ProA		.5	-	.5	-	.5	.5	.9	.05	.05	.5	.25	.25	.3	.4	.3	.4	.3	.3
M7 ProB		.5	-	.5	-	.5	.5	.5	.25	.25	.9	.05	.05	.05	.9	.05	.9	.05	.05

Table 2: Results of simulations. Harmony (i.e., degree of coherence of beliefs) after each new piece of evidence and distribution of finally preferred hypotheses

Distribution	Harmony									Preferred Hypothesis P(Hypothesis A)		
	ProA1	ProA2	ProA3	Amb1	Amb2	Amb3	ProB1	ProB2	ProB3	Overall	Strong Decay	Weak Decay
M1 fixed	.26	.39	.50	.46	.43	.41	.32	.25	.36	.80	0	1.0
M2	.25	.36	.46	.42	.40	.39	.31	.28	.33	.83	.43	.99
M3 ProA	.25	.36	.46	.46	.47	.49	.47	.46	.46	1.0	.99	1.0
M4 ProB	.25	.36	.46	.39	.35	.33	.29	.30	.35	.83	.49	.99
M5	.18	.23	.29	.27	.27	.28	.27	.28	.31	.83	.62	.92
M6 ProA	.18	.23	.29	.30	.33	.35	.34	.33	.34	.96	.90	.99
M7 ProB	.18	.23	.29	.24	.22	.22	.23	.32	.43	.36	.12	.58

The first distribution (M1) was identical to standard models and assumed no uncertainty about the coherence between evidence and hypotheses. The second (M2) closely resembled the standard model and assumed the same relations with a high probability of .9. The third distribution (M3ProA) represents shift of assumptions in favor of Hypothesis A after the first three pieces of evidence. The ambiguous evidence (Amb1-Amb3) is considered less supportive of Hypothesis B, and the evidence clearly favoring Hypothesis B (ProB1-ProB3) as less contradictory for A and less supportive of B. The forth distribution (M4ProB) represents a shift in favor of Hypothesis B. Now the ambiguous evidence is considered less supportive for Hypothesis A. If the model adequately captures the predictions of consistency theories, we should see an increase in coherence for M3 over M2 and M4.

The fifth probability distribution over coherence relations (M5) represents another set of basic assumptions. It assumes that all pieces of evidence are considered moderately supportive of the respective hypotheses. Distribution M6ProA again represents a shift of distribution M5 in favor of Hypothesis A while M7ProB represents a shift of M5 in favor of Hypothesis B. A comparison of the results for these distributions will show whether any of these shifts would increase coherence.

Model parameters were set to random values or kept constant for all simulations. Links of coherence had a weight of +.05, incoherence links of -.05. The incoherence link between hypotheses was set to -.2. Initial activations of hypotheses were set to random values between -.2 and +.2. Evidence nodes were added sequentially to the network after activations settled. Resulting activations were transferred to the next step. External activations received by evidence nodes were decayed when new evidence arrived. The decay parameter Λ was randomly set to values between 1 (strong exponential decay) and .1 (almost no decay). The activation added through the evidence nodes was kept constant at .5 for all steps. In line with previous studies we found that the qualitative pattern of activations hardly depended on the specific parameters (Thagard, 1989). Therefore only one set of results is reported here.

Results

For each probability distribution 10.000 constraint satisfaction models were instantiated and run. The results of the simulations are depicted in Table 2. Harmony, i.e., the resulting overall coherence of the belief network, is shown for each new piece of evidence. In addition, the percentage of cases in which Hypothesis A was preferred over B is given. For six out of seven distributions, Hypothesis A was preferred over Hypothesis B. Thus a primacy effect resulted, which is in accordance to the results of Kostopoulou et al., (2012). As expected, decay had a strong impact on results. When the decay was strong ($\Lambda=1$), that is, the last piece of evidence was strongly activated while previous evidence hardly received any activation, a recency effect sometimes occurred and Hypothesis B was favored. When decay was weak ($\Lambda=.1$), that is, initial evidence was activated only slightly less than the latest evidence, a primacy effect resulted even when the distributions were shifted in favor of Hypothesis B. Note that recent research indicates that weak or no decay fits best with people's actual decisions (Mehlhorn et al., 2011).

A comparison of distributions M1 and M2 shows that a probabilistic network with high probabilities basically results in the same overall preferences as a deterministic network which is identical to standard constraint satisfaction models. Overall coherence was only slightly reduced when relations became uncertain. A comparison of distributions M2, M3ProA and M4ProB indicates that the coherence of beliefs increased substantially when the probability distribution over coherence relations was shifted in favor of Hypothesis A, but not when it was distorted in favor of B. Note that an increase in coherence for M3ProA already resulted for the ambiguous items of evidence, after which it stayed at an elevated level. Thus the model predicts information distortion especially for the ambiguous items of evidence. This is what has been found empirically.

A comparison of distribution M5 to distributions M6ProA and M7ProB shows a different picture. Starting from less assertive assumptions about the diagnostic validity of the evidence, more coherence could be gained by shifting

assumptions in favor of Hypothesis B. A closer analysis shows that coherence increased for the ambiguous pieces of evidence by shifting assumptions towards Hypothesis A, but that this gain evaporated when the evidence favoring Hypothesis B arrived. Interestingly, a shift towards B only yielded substantially more coherence for the last few items. Thus the model predicts that people being uncertain should be less likely to distort but more likely to end up choosing Hypothesis B. This is what Kostopoulou and colleagues (2012) found.

General Discussion

A probabilistic constraint satisfaction model has been proposed to explain information distortion in sequential diagnostic reasoning. The model takes into account that diagnosticians may be uncertain whether a certain piece of evidence supports a diagnostic hypothesis for a particular case. To be more precise, it takes into account that people are aware of the fact that a piece of evidence may not always be present when a diagnosis is given and vice versa. Note that the model like constraint satisfaction models in general does not differentiate between the sensitivity of a diagnostic sign (i.e., the probability of the sign given the diagnosis) and the positive predictive value of the sign (i.e., the probability of the diagnosis given the sign). The model does, however, differentiate between believing a certain piece of evidence and believing that the information has diagnostic implications for a hypothesis.

The model has been implemented by using the standard formalism of ECHO (Thagard, 1989) and an exponential activation decay function to account for the sequentially arriving evidence. Uncertainty about diagnostic relations is represented by probability distributions over coherence relations among evidence and hypotheses. Belief in the observed evidence and hypotheses is represented by activations of the respective nodes.

Simulations were run to investigate the properties of the model and to find out whether it is able to predict findings reported in the literature. An analysis of the predictions of different probability distributions yielded several interesting results. First, the model shows a primacy effect which is reported frequently in the literature when people first receive several pieces of evidence favoring one hypothesis over others (Brownstein, 2003; Hogarth & Einhorn, 1992). However, many other models predict order effects, so this prediction is not unique.

Second, the model predicts information distortion. The results show that by distorting assumptions about the diagnostic validity of the observed cues, i.e., by shifting probability distributions over coherence relations, more coherent beliefs can be achieved. Importantly, coherence can be increased without giving up general assumptions about the coherence between cues and hypotheses. Thus, the model explains how the need for coherence can drive changes in beliefs about diagnostic validity and why information distortion may result.

Crucially, the simulations also showed that not all changes in beliefs about diagnostic relations may result in higher coherence. They also indicated that a shift in beliefs has an impact on coherence at a particular point during the diagnostic process. Thus the model allows for very specific predictions once initial beliefs about diagnostic relations are known.

Alternative Models

A number of parallel constraint satisfaction models has been proposed in the literature to account for diagnostic reasoning (e.g., Mehlhorn & Jahn, 2009; Mehlhorn et al., 2011; Wang et al., 2006; Gloeckner & Betsch, 2008; Gloeckner, Betsch, & Schindler, 2009).

The parallel constraint satisfaction model of Gloeckner and colleagues (2009) was designed to account for distortions in validity in multiple-cue judgment. In their research they found that participants changed their assessments of diagnostic validity depending on the favored option for a particular case. Although the model was devised for concurrent, non-sequential decision making it may be extended to cover sequential decision making. The structure of the model is highly similar to the model depicted in Figure 1 with cues being related to two alternative options, which compete with each other. Cues are assumed to be related to an activation unit. But, relations and activations are given an interpretation that is very different from our proposed model. The relation to the activation unit is assumed to represent the general validity of the cue, while the activation of each cue is considered to represent the validity of this cue for the particular case. This model is able to account for many findings in the judgment and decision making literature (cf. Gloeckner et al., 2009). Nevertheless, the model has difficulty to account for information distortion, because it does not differentiate between the validity of a cue and the diagnostic validity of the cue for a particular hypothesis. The results on information distortion (Kostopoulou et al., 2012) show that participants may increase the diagnostic validity with respect to Hypothesis A while decreasing the diagnostic validity with respect to Hypothesis B. The activation of a node, however, cannot increase and decrease at the same time. Thus the activation may represent whether a piece of evidence is considered valid or invalid, but not whether it is considered valid with respect to a diagnosis. The probabilistic constraint satisfaction model allows for this differentiation. Assumptions with respect to diagnostic validity are represented by probability distributions over coherence links. Therefore assumptions about the diagnostic validity with respect to several hypotheses may change independently from each other. Such a probabilistic constraint satisfaction model may account for the findings of Gloeckner et al. (2009). It also predicts that participants would lower their belief in the validity of cues contradicting the preferred option, as the resulting activation for these nodes would be negative.

The constraint satisfaction model proposed by Wang and colleagues UECHO (2006) was specifically developed to capture sequential belief updating and learning with a parallel constraint satisfaction network. The structure of the model is the same as the model shown in Figure 1. As outlined above, the model accounts for sequentially arriving evidence by a decay function over the activation distributed by the special activation unit. We adopted this idea for our model. The second important novel idea of UECHO is that the strength of the coherence links may change due to learning from feedback. We did not consider this idea for two reasons. First, clinicians are very unlikely to change their generic diagnostic knowledge in experimental studies on diagnostic reasoning and information distortion. Second, assigning specific weights to coherence links violates the fundamentally qualitative notion of coherence stressed by Thagard (1989). Either some evidence is coherent or incoherent with a hypothesis, or it is irrelevant. Our probabilistic model keeps this notion by assuming that links are either positive, negative or zero, while at the same time defining a probability distribution over these links representing the idea that evidence may be found even when the coherent hypothesis turns out to be false. Learning from feedback could be added to our model by adding a Bayesian learning algorithm that updates the probability distribution over coherence links. This seems to be a viable and elegant alternative to the proposal of Wang and colleagues (2006). In principle, UECHO may be extended to account for information distortion by assuming that the weight of individual coherence links may change for a particular case (i.e., change without learning). Like we envisioned for our probabilistic model, this shift may be driven by increased overall coherence (i.e., harmony). Such a model, however, would not be able to represent the uncertainty about coherence relations like the probabilistic model does.

Future Directions

The proposed probabilistic constraint satisfaction model shares important features with other parallel constraint models. In contrast to the other models it explains information distortion. As outlined above, the model allows for a number of specific predictions including conditions under which information distortion should not be found. However, to test these predictions, assumptions about diagnostic relations have to be assessed on an individual level and compared to later measurements of information distortion. Respective research still needs to be done. Hence, only future studies will show whether probabilistic constraint satisfaction models are able to successfully predict preferred hypotheses, information distortion and validity judgments of individual diagnosticians.

References

Brownstein, A.L. (2003). Biased predecision processing. *Psychological Bulletin, 129*(4), 545–68.

Chater, N. & Oaksford, M. (2008). *The probabilistic mind. Prospects for Bayesian cognitive science*. Oxford: Oxford University Press.

Croskerry, P. (2003). The importance of cognitive errors in medicine and strategies to minimize them. *Academic Medicine, 78*, 775–780.

Glöckner, A., & Betsch, T. (2008). Multiple-reason decision making based on automatic processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 1055–1075.

Gloeckner, A., Betsch, T., & Schlindler, N. (2009). Coherence shifts in probabilistic inference tasks. *Journal of Behavioural Decision Making, 23*, 439–462.

Hogarth, R. & Einhorn, H. (1992). Order effects in belief updating: the belief-adjustment model. *Cognitive Psychology, 24*, 1–55.

Kostopoulou, O., Delaney, B.C., & Munro, C.W. (2008). Diagnostic difficulty and error in primary care – a systematic review. *Family Practice, 25*, 400–413.

Kostopoulou, O., Russo, J.E., Keenan, G., Delaney, B.C. & Douiri, A. (2012). Information distortion in physicians' diagnostic judgments. *Medical Decision Making*. Published online. DOI: 10.1177/0272989X12447241

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive model of context effects in letter perception. Part 1. An account of basic findings. *Psychological Review, 88*, 375–407.

Mehlhorn, K., Taatgen, N.A., Lebriere, C., & Krems, J. (2011). Memory activation and the availability of explanations in sequential diagnostic reasoning. *Journal of Experimental Psychology: Learning, Memory & Cognition, 37*, 1391–1411.

Russo, J.E., Meloy, M.G., Wilks, T.J. (2000). Predecisional distortion of information by auditors and salespersons. *Management Science, 46*(1), 13–27.

Russo, J.E., Medvec, V.H., & Meloy, M.G. (1996). The distortion of information during decisions. *Organizational Behavior and Human Decision Processes, 66*, 102–110.

Simon, D., Snow, C. J., & Read, S. J. (2004). The redux of cognitive consistency theories: Evidence judgments by constraint satisfaction. *Journal of Personality and Social Psychology, 86*, 814–837.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences, 12*, 435–502.

Wang, H., Johnson, T.R., & Zhang, J. (2006). The order effect in human abductive reasoning: an empirical and computational study. *Journal of Experimental and Theoretical Artificial Intelligence, 18*, 215–247.