Structural Determinants of Counterfactual Reasoning

Daniel Yarlett (dany@cogsci.ed.ac.uk)
School of Cognitive Science, University of Edinburgh
2 Buccleuch Place, Edinburgh EH8 9LW Scotland

Michael Ramscar (michael@dai.ed.ac.uk)
School of Cognitive Science, University of Edinburgh
2 Buccleuch Place, Edinburgh EH8 9LW Scotland

Abstract

In this paper we explore the hypothesis that the processes underlying the matching and use of knowledge during counterfactual reasoning are the same as those that underlie reasoning by analogy. We report two experiments which indicate that, in common with analogical reasoning, counterfactual reasoning exhibits the following properties: (i) it is sensitive to systematic structural congruencies between representations; (ii) when background knowledge has to be retrieved before it can be exploited, the impact of structurally congruent background knowledge on the counterfactual inferencing process is mediated by featural commonalities.

Introduction

If it had not been for a young student at the University of Chicago in 1960, the Cuban missile crisis might have escalated to a nuclear war. At first blush this counterfactual sounds absurd, a *non sequitur* – how can there possibly be a link between a single student and the avoidance of nuclear war? It turns out, though, that the assertion isn't quite as far fetched as one might initially think:

> "Not yet twenty-one and too young to vote, the student worked in the Kennedy campaign. He was asked by the local Democratic organization if he would vote on behalf of a dead voter whose name was still on the rolls. He readily agreed and, refusing the small remuneration that was offered, forged the dead voter's signature and voted a straight Democratic ticket. The Illinois vote was close – Kennedy took the state by fewer than 10,000 votes – and critical. Illinois gave Kennedy the necessary electoral votes to win the presidency." (Lebow and Stein, 1996, p.119).

This additional knowledge connecting the student's actions to Kennedy's electoral success allows the rest of the story to fall into place. Had the student – and, crucially, others like him – not forged votes, Nixon may have been elected instead of Kennedy; if Nixon had been elected he would have established a much less liberal administration; a less liberal administration would have responded more forcefully to Khruschev's deployment of missiles in Cuba; and the odds are that this would have led to further military escalation and, ultimately, nuclear war.

This is not a watertight argument of course – the strength of some of the links in the argumentative chain is questionable – but it does serve to make apparent the relation between the two events which at first appeared utterly disparate.[1] And understanding that there is a connection between the existence of a group of students and the avoidance of nuclear war makes the initial counterfactual seem much more plausible than it did before: if those students hadn't existed then we can now see that this might just have tipped the balance in the favour of nuclear war. Background knowledge is clearly crucial to our assessment of counterfactuals, then, but how it is matched and used during counterfactual reasoning?

In this paper we explore the possibility that the processes underlying the matching and use of knowledge during counterfactual reasoning are the same as those that underlie reasoning by analogy. Although there has been extensive research into counterfactual reasoning over the last two decades from a cognitive-functional perspective (for representative overviews see Roese, 1997; Kahneman & Miller, 1986; Byrne & Tasso, 1999), no-one has explicitly investigated the effect of structural congruency on the inferencing process. We report two novel findings in support of the hypothesis that analogical mechanisms do underpin counterfactual reasoning in at least some contexts.

In Experiment 1 we show that an analogical match between background knowledge and a scenario about which subjects have to reason counterfactually significantly boosts their assessments of the soundness of related counterfactual inferences. This suggests that when engaged in counterfactual reasoning subjects rely on systematic structural matches between representations just as they do when making analogical inferences (Gentner, 1983; Holyoak and Thagard, 1995). Furthermore, previous studies in analogical inference and problem solving have found that it is difficult to exploit the information contained in analogies during real-life problem-solving because analogies are typically difficult to retrieve (Gick and Holyoak, 1980; Gentner, Ratterman and Forbus, 1993). In Experiment 2 we add a retrieval requirement to the task investigated in Experiment 1, and show that this same finding also applies to counterfactual reasoning.

Why Analogy?

Theories of counterfactual reasoning agree that similarity plays a core role in determining legitimate inferences – in order to evaluate what would be true if A were the case, it makes sense to consider only the state of affairs that is most similar to the way things actually are, except that in

[1] See Lebow and Stein (1996) for a fuller analysis of the plausibility of this counterfactual sequence of events.

it *A* is true instead of false. This similarity constraint prevents unnecessary or extraneous alterations to the situation under consideration that could otherwise bias the inferencing process, and is what makes counterfactual reasoning distinct from the idle entertainment of hypothetical situations. For example, consider the relatively uncontentious counterfactual 'If this match had been struck, it would have lit'. As it stands this is most likely a true thing to say (depending on the precise context of utterance), but without the similarity constraint one can trivially refute it by making unwarranted changes to the basic situation and hence considering situations where matches are non-flammable, or where all matches are underwater, and so on.

However, despite the central role that similarity plays in theories of counterfactual reasoning, no consensus exists on how it should be defined. Possible world theorists treat it as a primitive partial ordering between possible worlds (Lewis, 1973; Stalnaker, 1968); law-based theorists treat it as the preservation of consequences of natural laws (Chisholm, 1946; Goodman, 1947; Pollock, 1976); whilst others treat it as minimal adjustments to autonomous causal mechanisms represented in dependency structures (Pearl, 2000). These formal characterisations of similarity and the constructs they rely on make no claim to being cognitively plausible, and are furthermore open to charges of underspecification (Lewis, 1973, §4.2; Pollock, 1976, p.17; Fine, 1975; and Bowie, 1979).

In contrast, work on the notion of similarity in cognitive science has concentrated on the empirical testing of concrete proposals about what it means for two things to be similar to one another. One finding is that the shared and distinct features of two representations play an important role in determining their similarity (Tversky, 1977), but more recent work has also shown that the *relational structure* holding between features also affects judgements of similarity (Gentner & Markman, 1997; Love, 2000).

Gentner and Markman (1997) point out the effect of structure on similarity judgements is consistent with the idea that analogical alignment plays a role in the evaluation of similarity. If analogy really does play a role in determining similarity then, given that similarity is so central to counterfactual reasoning, it makes sense to investigate whether analogy plays a role in counterfactual reasoning by testing for effects of structure on the inferencing process. Furthermore, the account of similarity provided by the analogy-based *structural alignment* framework (Gentner, 1983) has several merits from a cognitive perspective.

Firstly, the structural alignment framework has been subject to a program of empirical testing which provides evidence regarding its plausibility as an account of the way that people compare representations (see Gentner & Markman, 1997, for a review). Secondly, the structural alignment process can be modelled computationally (e.g. SME; see Falkenhainer, Forbus and Gentner, 1989) meaning that the processes it relies upon are tractable. The final reason for preferring the structural alignment account of similarity over alternative accounts is that the analogical mechanisms that it is founded upon have been explicitly connected with many other cognitive skills – most relevantly inference (see Holyoak and Thagard, 1995, for a good review of the many applications of analogy) – offering yet further reason to believe that it may also be implicated in the process of counterfactual reasoning at the cognitive level.

	+ Structure	− Structure
+ Features	Literal Similarity	Surface Similarity
− Features	Analogy	First-Order Relations

Table 1: The commonalities each variant category shares with its base in the materials used by Gentner, Ratterman & Forbus (1993).

Given these reasons, the question of whether the same empirical findings about analogical inference can also be identified in counterfactual reasoning tasks is a pertinent one. If commonalities between counterfactual and analogical reasoning could be established empirically, this would undoubtedly be profitable to our understanding of counterfactuals due to the mature status of research into theories and models of analogy.

Previous Work on Analogical Matching

Experiments conducted by Gentner, Ratterman and Forbus (1993) investigated the effect that systematically varying featural and structural matches between two representations of knowledge had on three measures: the *retrievability* of one item given the other as a prime; on assessments of *inferential soundness*[2] between the two items; and on the perceived *similarity* of the two items. Their series of studies used materials consisting of a *base* scenario and four *variants* of this base in the 2×2 design depicted in Table 1. The systematic variation of featural and structural matches in this design allowed the influence of both factors on a range of tasks to be assessed, and results in a taxonomy of four different types of variants to a base: literal similarity (LS), surface similarity (SS), analogy (AN) and first-order relations (FO).

An important finding reported in this series of studies was that the inferential soundness of two domains is determined primarily by the sort of systematic structural congruencies that analogy is sensitive to. Experiment 1 tests whether this result applies to counterfactual reasoning too, by measuring whether analogical matches to background knowledge boost evaluations of the soundness of counterfactual inferences in the same way as they do in analogical reasoning.

An additional finding about analogical inference is that structural matches can only be exploited in the inferencing process if they are first retrieved, which is typically a difficult thing for people to do (Gick and Holyoak, 1980). This suggests another hypothesis that can be tested: if the cognitive processes underlying counterfactual reasoning are really like those underlying analogy, then in contexts that demand retrieval we would expect analogy based inferences to occur relatively infrequently, because the impact of structural matches on the counterfactual inferencing process is mediated by the featural similarities that drive retrieval (Gentner, Ratterman and Forbus, 1993; Ramscar and Yarlett, 2000). This prediction is tested in Experiment 2.

[2]Inferential soundness is defined as the degree to which knowledge about one domain can be used to generate appropriate inferences about another domain.

Experiment 1

Analogy is universally held to be determined by systematic congruencies in the relational structure of two representations. If analogical mechanisms really do underpin counterfactual reasoning then we would expect to find that structural matches support counterfactual reasoning in exactly the same way that they support analogically based reasoning (Falkenhainer, Forbus & Gentner, 1989; Gentner, Ratterman & Forbus, 1993). Moreover, we would expect a simple featural match between a counterfactual reasoning problem and background knowledge to have no such facilitatory effect.

The simplest test for an effect of shared structure as opposed to shared features in counterfactual inference, as predicted by the hypothesis that analogical mechanisms underpin counterfactual reasoning, is to place subjects in two conditions: one in which they have information that counterfactually matches the problem scenario structurally but not featurally (AN condition); and the other in which the opposite is the case (SS condition). The effect of the knowledge on the soundness ratings assigned a counterfactual inference related to the two sets of background knowledge can then be measured relative to a pre-test condition (the *a priori* or AP condition) in which subjects are provided with *no* background knowledge relevant to the counterfactual in question. If counterfactual reasoning really is sensitive to analogical matches then we would expect a significant increase from the pre-test soundness ratings in the AN condition, but no such increase in the SS condition.

Method

Subjects read either the SS or AN match to a base scenario, and were then asked two comprehension questions about it to ensure they had read it properly. They were then asked to read the corresponding base, and reason counterfactually about it: subjects had to rate the soundness of four counterfactual inferences which undid a key event described in the base. Both SS and AN variants were designed to promote the same counterfactual inference by featuring it as the outcome of the sequence of events they described. The relative effect of a featural versus a structural match between the base problem and background knowledge on the inferencing task could therefore be measured by comparing the soundness ratings for the same inference in each condition. Soundness ratings were elicited by asking subjects how likely the counterfactual inference was to be true. Ratings were provided on a 9 point scale anchored by 'extremely unlikely' at the lower end, and 'extremely likely' at the upper end.

Participants Participants were 40 undergraduate students of the University of Edinburgh, completing an introductory course in psychology. All participants were volunteers, and no reward was offered for taking part in the experiment.

Materials The materials were designed to replicate the experimental design used in Gentner, Ratterman and Forbus (1993) as summarised in Table 1, while also being compatible with a counterfactual reasoning task. 5 base descriptions and 4 variant categories for each base (corresponding to the LS, SS, AN, FO categories described earlier) were designed, resulting in 25 scenarios in total. Only the SS and AN variants were used in Experiment 1.

> Ostavia, Grern and Donnol were three neighbouring and hostile states that seemed to take a particular delight in antagonising one another. It wasn't uncommon these days to hear of one country criticising the others, and threatening them with military action. The trouble had begun several decades ago when Ostavia had attacked both Grern and Donnol in a failed attempt to invade them. Since then relations between the three states had gone downhill, and just recently things appeared to be getting even worse.
>
> **Ostavia had recently launched an attack against Grern**, and although resistance was fierce it looked as though Ostavia was soon going to conquer the Grern forces. However, Grern approached Donnol for help in the conflict, arguing that if they didn't prevent the onslaught then Ostavia would be coming for them next. Donnol considered Grern's argument and saw that its conclusion was probably correct. Therefore Donnol sent its troops to reinforce Grern's units, and after an intense struggle the Ostavian attack was repelled. *Soon after, Grern and Donnol formed a long-lasting alliance, and occupied Ostavia enforcing harsh sanctions on its citizens.*

Figure 1: Example base description. The key event for this description is shown in bold type, and the outcome altered in the variant conditions in italics (distinctions were unmarked in experimental materials).

Each base was structured so that it described a key event which led to a specific outcome (see Figure 1 for an example). For each base, four alternative counterfactual outcomes to the actual outcome described in the base were selected (these were the alternatives rated for soundness by subjects). The counterfactual analogy (AN) variant to the base in Figure 1 is shown in Figure 2. The counterfactual inferences corresponding to the example base scenario that subjects had to rate for soundness are also shown, in Figure 3.

Each of the four variants of a particular base was designed to undo the key event described in the base while either sharing or not sharing featural and structural commonalities with it, depending on which of the variant categories it belonged to. All variants in a material set were constructed to end in an identical outcome, different from that occurring in the base of the set (in the example materials shown the four variants ended with some form of reconciliation between the parties involved, instead of the 'harsh sanctions' mentioned in the base). The four variants can thus be regarded as providing different types of background knowledge, each of which could serve to promote the plausibility of the same counterfactual inference. The issue of interest in this paper is how these four categories of background knowledge, and the different ways that they are related to the base scenario, affect subjects' soundness judgements about the counterfactual outcomes they describe.

The four variants to a base were created in the following manner. First, a key event in each base was identified. In those variants to a base in which a counterpart to the key event should be easily discernible (i.e. in the LS and AN structurally congruent variants), the appropriate counterpart was explicitly negated to make it counterfactual to the base with respect to the key event. In those variants in which a counterpart of the key event should not be easily identifiable

A chess-playing craze had recently swept through the inhabitants of Chesterton, and as a result three chess clubs had sprung up in the sleepy town over the last few years. The clubs were, unimaginatively enough, called the Maters, the Gambit Players and the Rank and Filers. Times were hard for the clubs, because there simply weren't enough players to go around. The problems had been started by the antagonistic members of the Maters. Several months earlier the Maters had scandalously offered reduced subscription rates to members of the Gambit Players and Rank and Filers. Since then an air of hostility had descended upon the chess-playing community in Chesterton.

The Maters thought about trying again to gain extra members from the other two clubs, this time by promising them extra facilities as well as reduced subscription rates. **However, eventually the Maters thought this was a bad idea because it seemed too underhand and risky.** Unbeknownst to the Maters, the other two clubs had formed a pact with one another that they would merge and drive the Maters out of town if they tried to steal any of their members again. *However, because the Maters thought better of their plan everyone relaxed a little, and a couple of years later the clubs were on very friendly terms.*

Figure 2: Example analogy (AN) variant. The counterfactual counterpart of the key event is shown in bold type, and the counterfactual outcome is shown in italics (distinctions were unmarked in experimental materials).

If Ostavia hadn't decided to launch an attack against Grern and Donnol, it wouldn't have been defeated in war *and something else would probably have happened*. Rate the following outcomes according to how likely you think they would have been to occur if Ostavia hadn't launched its attack (please circle *one* number for each of the four options):

1. **All three countries would have ended up disarming, and developing friendly relations.**
2. The countries would have organised a yearly carnival in celebration of newfound connections.
3. The countries would have merged under an even greater threat.
4. The countries would not have learnt from their mistakes, and would end up fighting one another again.

Figure 3: Example counterfactual outcomes that subjects were asked to rate for soundness. The inference promoted by the base's four variants is shown in bold type.

(i.e. in the SS and FO conditions, where there is no structural congruity with the base), no event was described that could be construed as being an instance of the key event.

Procedure Subjects were exposed to four of the material sets, two in the SS condition and two in the AN condition. Materials were randomised across subjects, and order of presentation was counterbalanced.

The written instructions to the experiment informed participants that they were taking part in a study investigating general reasoning, and that they were going to be shown a series of scenarios and asked to make some inferences about them. Participants were further told that for each problem they were going to be provided with some background information (this information corresponding to the SS and AN conditions of the experiment), and then asked to reason about a problem scenario. They were explicitly informed that the background information **may** be of use, and that they should only make use of it if they thought it was relevant to the problem in hand.

Pre-test

A pre-test condition was run before Experiment 1 to allow the facilitatory effects of providing various categories of background knowledge to be compared to the case in which subjects are provided with *no* relevant background knowledge. 20 postgraduate students of the Division of Informatics, University of Edinburgh, were provided with the 5 base descriptions in a semi-randomised order, and asked to evaluate how sound it seemed that each of 4 counterfactual alternatives would occur had the key event in each description been undone. Ratings were elicited on a 9 point scale. Two sets of booklets presented all 5 base descriptions and the corresponding counterfactual inferences to be rated in reverse orders to minimise order effects. Subjects were randomly assigned to either of the two booklet conditions.

Results

The mean soundness ratings of the promoted inferences in the SS- and AN-exposure conditions from Experiment 1, along with the pre-test (AP) soundness ratings for the same inferences, are shown in Figure 4. T-testing revealed no significant shift in soundness ratings from the pre-test to the SS condition (Welch's $t = 0.36$, $df = 58.05$, $p > 0.025$, one-tailed). In contrast, soundness ratings did increase significantly from the pre-test to the AN condition (Welch's $t = 2.00$, $df = 58.64$, $p < 0.025$, one-tailed).[3]

Discussion

The results of this experiment indicate that subjects will reason counterfactually on the basis of an analogical match to background knowledge when one is available, and that a featural match to background knowledge alone is insufficient to support a counterfactual inference. The fact that an analogical match to background knowledge makes a counterfactual inference seem more sound shows that people are sensitive to the sort of systematic structural correspondences found in analogy when engaged in this form of reasoning, and therefore supports the hypothesis that analogical mechanisms underpin counterfactual reasoning at the cognitive level.

Experiment 2

Gentner, Ratterman and Forbus (1993) present evidence that in contexts where subjects first had to retrieve knowledge in order to be able to exploit it to make an analogical inference

[3]Both tests involved samples with unequal variances by Levene's test, and hence Welch's corrected t statistic was calculated. Tests were conducted at adjusted α levels to keep the overall risk of a Type I error at the 0.05 level, by solving $1-(1-\alpha)^2 = 0.05$ for α.

(i.e. where relevant knowledge is not presented in association with the base it pertains to as it was in Experiment 1), both structural *and* featural matches play a role in determining the usefulness of that knowledge in the inferencing process. The effect of featural commonalities is introduced because the process of retrieval is sensitive to featural, but not structural, matches (Gentner, Ratterman and Forbus, 1993; Ramscar and Yarlett, 2000). Therefore, if counterfactual reasoning really is similar to analogy at the cognitive level, we would expect the addition of a retrieval requirement to the inferencing task of Experiment 1 to have the following effects.

First, the AN condition should no longer produce a significant boost in soundness ratings compared to the pre-test condition. We expect this because AN matches are typically difficult to retrieve (see, for example, Gick and Holyoak, 1980), and are therefore less likely to be used in the counterfactual inferencing process in this context. Second, the soundness ratings of SS condition inferences should remain the same as in Experiment 1 (because SS materials have high retrievability the additional retrieval requirement should make no difference to the availability of this knowledge in the inferencing process). Third, and finally, the overall pattern of results over the 4 variant categories should exhibit sensitivity to both featural and structural factors, instead of just structural factors as in Experiment 1 (because featural matches facilitate retrieval, which mediates the influence of structure on the inferencing process).

Experiment 2 was accordingly designed to test these hypotheses by adding a retrieval task to Experiment 1, so that background knowledge had to be successfully retrieved before it could be used as the basis for supporting a counterfactual inference.

Method

In the first phase of the experiment subjects were asked to read 5 variants, one from each material set. Two multiple-choice comprehension questions were asked after each story had been read, in an attempt to ensure that subjects read the information thoroughly. They were next asked to complete one of two short distraction tasks. The distraction tasks were unrelated to the present study; one investigated a spatio-temporal priming phenomenon, and the other morphological inflection. In the second phase of the experiment subjects were asked to perform exactly the same inferencing task as described in Experiment 1, except that this time they had also been provided with the LS and FO variants of background knowledge, in addition to the SS and AN categories. However, because the connection between the base problems and the background knowledge they were provided with in the variants was not made explicit, subjects had to retrieve relevant knowledge before being able to use it in the inferencing process. Subjects were not told that there was any connection between the first and second part of the experiment. This was done to determine whether subjects would be able to *spontaneously* recruit (Kahneman and Miller, 1986) background knowledge in order to make their soundness judgements for the counterfactuals.

Participants Participants in this study were 148 undergraduates and postgraduates at the University of Edinburgh. All participants were volunteers. The undergraduate students were awarded class credit for taking part in the study.

Materials All 5 bases, with their corresponding 4 variants (LS, SS, AN and FO categories), of the materials described in Experiment 1 were used.

Procedure 1 variant was taken from each of the 5 material sets, to result in each prime set. This meant that there were 4 prime sets in total. Two random orderings of each prime set were combined with two random orderings of the question set (which was the same as that used in the pre-test and Experiment 1), to create four distinct material orderings for each of the four prime sets. This was done to minimise the potential for presentation effects, and resulted in 16 types of booklet in total. The 2 distractor tasks were randomly included in the booklets, and took subjects less than five minutes to complete.

In the first part of the experiment subjects were instructed only that they were taking part in a memory experiment, and that they should therefore read the stories they were presented with carefully. In the second part of the experiment no reference was made to the stories in the first part; subjects were merely asked to provide soundness ratings as they saw fit.

Results

The results of Experiment 2 are shown in Figure 4 (along with the results from Experiment 1).

A 2×2 ANOVA was conducted on the four experimental conditions. This revealed a significant effect of featural matches ($F(1,147) = 8.75, p < 0.01$), a significant effect of structural matches ($F(1,147) = 16.76, p < 0.01$), and no significant interaction between the two factors ($F(1,147) = 3.72, p > 0.05$) consistent with the predictions of the analogy-based account of counterfactual reasoning.

The soundness ratings in neither the SS (Welch's $t = 0.81$, $df = 49.03$, $p > 0.025$, one-tailed) nor the AN (Welch's $t = 0.90$, $df = 51.61$, $p > 0.025$, one-tailed) conditions differed significantly from those in the pre-test.[4] Whilst the effect of SS knowledge on the counterfactual inferencing process is unchanged between Experiments 1 and 2, the facilitation of soundness ratings produced by the analogical knowledge observed in Experiment 1 has disappeared. The best explanation of this difference is that the retrieval requirement introduced in Experiment 2 has prevented AN matches from being successfully retrieved in the counterfactual inferencing process, a finding paralleled in analogy research.

Discussion

The results of Experiment 2 were consistent with those predicted by the analogy-based account of counterfactual reasoning. Whereas the provision of analogical knowledge in Experiment 1 was sufficient to boost the soundness of related counterfactual inferences, the same effect of analogical knowledge was not observed in Experiment 2. The best explanation of this seems to be that retrieval does appear to mediate the effect of structural congruency in the solution of

[4] Corrected α values were used as before. Levene's test indicated both t-tests involved samples with unequal variances, and so Welch's corrected t was calculated.

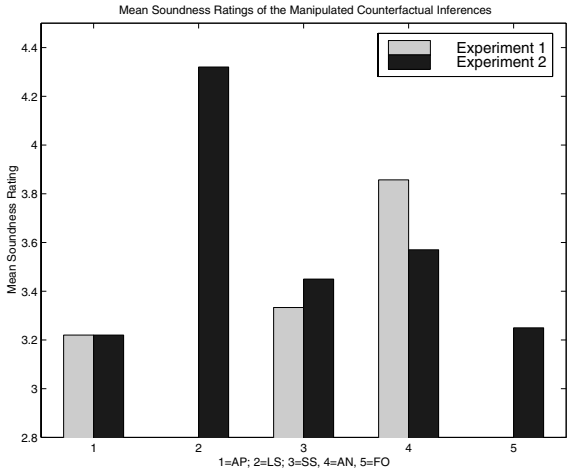

Figure 4: A bar chart plot of the mean plausibility of the manipulated counterfactual inferences in the pre-test and Experiments 1 and 2.

the counterfactual reasoning problems. Furthermore, both featural and structural factors were shown to have a significant effect on the soundness ratings provided by subjects, suggesting that both a feature sensitive retrieval process and structural alignment have a role to play in accounting for the way in which background knowledge is exploited during counterfactual reasoning. This pattern is in accordance with the research on analogical reasoning reported by Gentner, Ratterman and Forbus (1993).

General Discussion

We have presented evidence that there are deep similarities in the ways in which knowledge is exploited during both analogical and counterfactual reasoning. Specifically, we have shown that an analogical match to stored knowledge of an appropriate kind is sufficient to significantly boost the plausibility of a counterfactual inference, and that the structure-matching process that we hypothesise underpins counterfactual inference is mediated by a retrieval process sensitive to featural-matches, just as in analogical reasoning. The former finding is most significant, perhaps, because it is unpredicted by current theories of counterfactual reasoning.

Whilst there are limitations to the current work – chiefly that it doesn't explore what occurs when complete matches to stored counterfactual knowledge are unavailable, and counterfactual representations have to be *constructed* – we nevertheless feel that the current research presents the intriguing possibility of applying existing findings about to counterfactual reasoning, in order to improve our understanding of this important cognitive process.

Acknowledgements

We would like to thank Scott McDonald, Keith Stenning, Lera Boroditsky and the three anonymous reviewers for useful comments on the work reported in this paper.

References

Bowie G.L. (1979). The Similarity Approach to Counterfactuals: Some Problems. *Noûs*, 13, 477-498.

Byrne R.M.J. and Tasso A. (1999). Deductive Reasoning with Factual, Possible, and Counterfactual Conditionals. *Memory & Cognition*, 27(4), 726-740.

Chisholm R.M. (1946). The Contrary-to-Fact Conditional. *Mind*, 55(220), 289-307.

Falkenhainer B., Forbus K.D. and Gentner D. (1989). The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, 1-63.

Fine K. (1975). Critical Notice of *Counterfactuals*. *Mind*, 84, 451-458.

Gentner D. (1983). Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science*, 7, 155-170.

Gentner D., Ratterman M.J. and Forbus K.D. (1993). The Roles of Similarity in Transfer: Separating Retrievability from Inferential Soundness. *Cognitive Psychology*, 25, 524-575.

Gentner D. and Markman A.B. (1997). Structural Alignment in Analogy and Similarity. *American Psychologist*, 52(1), 45-56.

Gick and Holyoak (1980). Analogical Problem Solving. *Cognitive Psychology*, 12, 306-355.

Goodman N. (1947). The Problem of Counterfactual Conditionals. *The Journal of Philosophy*, 44(5), 113-128.

Holyoak K.J. and Thagard P. (1995). Mental Leaps: Analogy in Mental Thought. MIT Press, Cambridge, Massachusetts.

Kahneman D. and Miller D.T. (1986). Norm Theory: Comparing Reality to its Alternatives. *Psychological Review*, 93(2), 136-153.

Lebow R.N. and Stein J.G. (1996). Counterfactuals and the Cuban Missile Crisis. In Tetlock P.E. and Belkin A. (eds.), *Counterfactual Thought Experiments and World Politics: Logical, Methodological and Psychological Perspectives*, Chapter 5, p.119-148, Princeton University Press, Princeton, New Jersey.

Lewis D.K. (1973). Counterfactuals. Harvard University Press, Cambridge, Massachusetts.

Love B. (2000). A Computational Level Theory of Similarity. *Proceedings of the 22nd Annual Meeting of the Cognitive Science Conference*, 316-321.

Pearl, J. (2000). Causality: Models, Reasoning, and Inference. Cambridge University Press, Cambridge.

Pollock J.L. (1976). Subjunctive Reasoning. Volume 8 of *Philosophical Studies Series in Philosophy*, D. Reidel Publishing Company, Dordrecht, Holland.

Ramscar M.J.A. and Yarlett D.G. (2000). A High-Dimensional Model of Retrieval in Analogy and Similarity-Based Transfer. *Proceedings of the 22nd Annual Meeting of the Cognitive Science Conference*, 381-386.

Roese N.J. (1997). Counterfactual Thinking. *Psychological Bulletin*, 121, 133-148.

Stalnaker, R. (1968). A Theory of Conditionals. In Rescher, N. (ed.) Studies in logical theory, number 2 in American Philosophical Quarterly Monograph Series, Blackwell Press, Oxford.

Tversky, A. (1977). Features of Similarity. *Psychological Review*, 84(4), 327-352.