

# An Evaluation of ChatGPT-4’s Qualitative Spatial Reasoning Capabilities in RCC-8

A G Cohn

University of Leeds & The Alan Turing institute, UK

## Abstract

Qualitative Spatial Reasoning (QSR) is well explored area of Commonsense Reasoning and has multiple applications ranging from Geographical Information Systems to Robotics and Computer Vision. Recently many claims have been made for the capabilities of Large Language Models (LLMs). In this paper we investigate the extent to which one particular LLM can perform classical qualitative spatial reasoning tasks on the mereotopological calculus, RCC-8.

## Introduction

Qualitative Spatial Reasoning (QSR<sup>1</sup>) (Cohn and Renz 2008; Chen et al. 2015; Cohn and Hazarika 2001) is a well developed field which is concerned with the representation of qualitative spatial information and reasoning with it. In natural language, spatial information is usually represented qualitatively (using prepositions such as *on*, *in*, *left of*, *part of*, *under*, *touching*, ...) and many calculi have been developed to represent such information. There are calculi for mereological relations (such as RCC-5 (Jonsson and Drakengren 1997)), mereotopological relations (such as RCC-8 (Randell, Cui, and Cohn 1992; Cohn et al. 1997)), directions (such as OPRA (Moratz 2006)), size (Gerevini and Renz 2002) for example as well as calculi combining two different aspects of spatial information, such as the Rectangle Algebra (Guesgen 1989; Mukerjee and Joe 1990) which can represent both mereotopological information as well as directional. What is common to all these calculi is that they consist of a set of *jointly exhaustive and pairwise disjoint* (JEPD) *base* relations. For example, RCC-8 contains eight JEPD *base* relations, illustrated in 2D in Fig. 1.

*Large Language Models* (LLMs) (Devlin et al. 2019; Brown et al. 2020), such as ChatGPT-4 (Roumeliotis and Tselikas 2023) are a recent example of so called *Foundation Models* which have been trained on very large textual corpora in order to generate text in response to a prompt. This is not the place to survey this burgeoning field, but we note that many claims have been made for the power and apparent intelligent behaviour that these models can display. In particular their performance on some benchmarks may lead one

Copyright © 2023, The Authors. All rights reserved.

<sup>1</sup>We may use QSR as shorthand for both Qualitative Spatial Reasoning and Qualitative Spatial Representation; context should usually make clear which is intended.

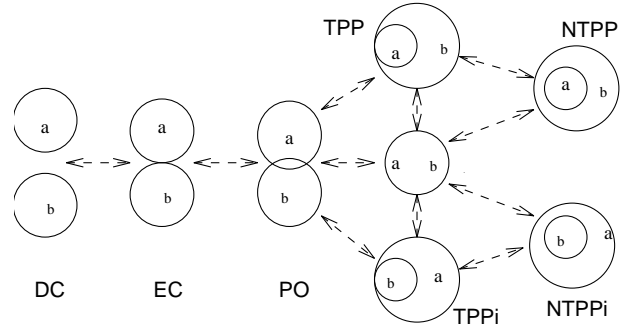


Figure 1: The eight relations of the RCC-8 calculus illustrated in 2D.

to believe that they possess, at least to some degree, the ability to perform commonsense reasoning. Spatial reasoning is usually regarded as one core aspect of common sense so it is natural to ask whether LLMs can reason about qualitative spatial information. This is the question that we address here.

In earlier work (Cohn and Hernandez-Orallo 2023) we use extended dialogues with an LLM to try to map the boundaries of spatial commonsense in some LLMs, addressing a variety spatial challenges, and examining not only the response given but also the explanation/justification of the response, but did not specifically focus on existing QSRs, though some questions were asked which do correspond to particular reasoning steps in an existing QSR. Here we focus on one specific QSR and ask the question as to what extent an LLM can perform reasoning in that calculus, and conduct a more exhaustive evaluation, but looking at the ability to perform compositions between relations and also to reason about the conceptual neighbourhood diagram of the calculus. Weaknesses in the reasoning powers of LLMs have previously been noted (e.g. (Cai, Chang, and Han 2023)) so one might not expect LLMs to perform well in this regard. But on the other hand, there are a large number of papers about QSR in the literature and these are likely to have formed part of the training corpus of an LLM, and thus might facilitate correctly responding to prompts – though the information concerning the actual reasoning steps are often given in tables (in particular *composition tables* – see below) and thus

might be hard for LLM training procedures to process well.

There are now many LLMs in the literature. Some of these are open source and are explicit about the training corpus; others are closed and give no specific information about the training, or the precise corpus, such as the GPT family of LLMs. Nevertheless since we observed previously (Cohn and Hernandez-Orallo 2023) that ChatGPT-4 and GPT4 were the most performant for spatial reasoning, we use ChatGPT-4 as the LLM with which we perform our experiments. In each of the experiments below, an initial prompt gave the problem setting and the task to be performed. Subsequent prompts in the conversation probed about one specific inference (e.g. one cell in a composition table). Each experiment was a separate conversation and the version number of ChatGPT-4 was given as “May 24” (presumably, the version of May 24 2023). The conversations can be found at the following location: <https://tinyurl.com/qr23sup>.

## Compositional Reasoning

The most researched form of reasoning with Qualitative Spatial Relations is that of composition: i.e. given a two facts  $R1(x, y)$ , and  $R2(y, z)$ , then what relations are possible between  $x$  and  $z$ ? In general more than one base relation is possible, so the answer may take the form of a disjunction of relations, as can be seen in the composition table for RCC-8 in Fig. 2. Given the ubiquity of RCC-8 in the QSR literature, in this paper we focus on the abilities of an LLM to reason with RCC-8.

### Experiment 1: Compositional Reasoning in RCC-8

In this experiment the standard names of the eight relations were given as well as a sentence stating that the relations were part of the RCC-8 calculus:

*Consider the following set of eight pairwise disjoint and mutually exhaustive binary spatial relations. These relations form part of the well known RCC-8 qualitative spatial reasoning calculus.  $DC(x,y)$  means that  $x$  and  $y$  are disconnected and share no spatial parts.  $EC(x,y)$  means that  $x$  and  $y$  touch at a boundary but do not share any interior parts.  $PO(x,y)$  means that  $x$  and  $y$  share a spatial part, but neither is part of the other.  $TPP(x,y)$  means that  $x$  is part of  $y$  and touches  $y$ 's boundary.  $NTPP(x,y)$  means that  $x$  is part of  $y$  but does not touch  $y$ 's boundary.  $TPPi(x,y)$  is the same as  $TPP(y,x)$ .  $NTPPi(x,y)$  is the same as  $NTPP(y,x)$ . Finally,  $EQ(x,y)$  means that  $x$  and  $y$  are coincident. I will now ask you a series of questions about these relations. There may be more than one possible relation, in which case name all of the possible answers. If all eight relations are possible, just say ALL.*

As an example of a prompt and ChatGPT-4's response, consider the first one:

**Prompt:** *If  $DC(x,y)$  and  $DC(y,z)$  then what are the possible relationships between  $x$  and  $z$ .* **Response:** *If  $DC(x,y)$  and  $DC(y,z)$ , then  $x$  and  $z$  can either be disconnected ( $DC$ ) or externally connected ( $EC$ ). This*

*is because the disconnection between  $x$  and  $y$ , and  $y$  and  $z$ , only guarantees that neither  $x$  nor  $z$  is connected to  $y$ . It does not preclude the possibility that  $x$  and  $z$  themselves might be disconnected or touching at a boundary. However, the other relations like part-of ( $TPP$ ,  $NTPP$ ) and their inverses ( $TPPi$ ,  $NTPPi$ ), partial overlap ( $PO$ ), and equality ( $EQ$ ) are not possible because they would imply some level of connection between  $x$  and  $z$  through  $y$ , which contradicts our initial conditions. So, the possible relationships between  $x$  and  $z$  are:  $DC(x,z)$ :  $x$  and  $z$  are disconnected.  $EC(x,z)$ :  $x$  and  $z$  are externally connected.*

Notice that in fact the composition of  $DC$  with  $DC$  is the union of all eight RCC-8 relations, so the reasoning is incorrect here.

The results of the  $7*7=49$  prompts issued (compositions involving  $EQ$  were omitted since they should be trivial) are summarised in Fig. 3. There is a total of 85 relations correctly predicted (those in black), 197 relations which are correctly not predicted (those not in the table). Thus there are a total of 282 correctly predicted (both present and missing) relations, i.e. 71.94% of the total possible ( $7*7*8 = 392$ ). 11 of the cells (out of 49) are completely correct. This is clearly much better than chance but far from a correct composition table. Fig. 4 shows the absolute number and relative percentages of the 8 relations in the proportion that they are correctly or incorrectly predicted. As can be seen,  $DC$  is never missed as a correct prediction, and  $PO$  is never incorrectly predicted as present.  $DC$  has the highest accuracy which is perhaps not surprising since  $DC$  might be regarded as intuitively easier to reason about (and this fits with the predictions about human preferred models in the section below). Including both correctly predicted and correctly missing relationships, all eight relations have broadly similar accuracies.

In order to test whether the result was influenced by prior knowledge of RCC-8 gained as part of its training, we also performed the same experiment, but with all the relation names prefixed by an X to disguise the connection to RCC-8. The prompt was the same as above except for the change of relation names and the omission of the second sentence. The results are given in Fig. 5 while Fig. 6 shows the absolute number and relative percentages of the 8 relations in the proportion that they are correctly or incorrectly predicted. As can be seen,  $DC$  again is never missed as a correct prediction, and  $EC$  is only missed twice; again  $PO$  is never incorrectly predicted as present. As before,  $DC$ ,  $EC$ , and  $PO$  have the highest accuracies, along with  $EQ$ , but  $EQ$  is never predicted as present correctly, only incorrectly. The overall average of correctly predicted relations (present and missing) drops from 71.94% in the non-anonymous case above to 67.09% so there is some loss of performance though whether is due to the anonymisation of the relations or the stochasticity of ChatGPT-4 is not clear.

### Experiment 2: Preferred Compositions in RCC-8

As noted above, in general a composition of two relations will yield more than one possible base relations, but it turns

|       | DC                      | EC                       | PO                    | TPP                   | NTPP                          | TPPi                     | NTPPi                   | =     |
|-------|-------------------------|--------------------------|-----------------------|-----------------------|-------------------------------|--------------------------|-------------------------|-------|
| DC    | *                       | DC, EC, PO, TPP, NTPPi   | DC, EC, PO, TPP, NTPP | DC, EC, PO, TPP, NTPP | DC, EC, PO, TPP, NTPP         | DC                       | DC                      | DC    |
| EC    | DC, EC, PO, TPPi, NTPPi | DC, EC, PO, TPP, TPPi, = | DC, EC, PO, TPP, NTPP | EC, PO, TPP, NTPP     | PO, TPP, NTPP                 | DC, EC                   | DC                      | EC    |
| PO    | DC, EC, PO, TPPi, NTPPi | DC, EC, PO, TPPi, NTPPi  | *                     | PO, TPP, NTPP         | PO, TPP, NTPP                 | DC, EC, PO, TPPi, NTPPi  | DC, EC, PO, TPPi, NTPPi | PO    |
| TPP   | DC                      | DC, EC                   | DC, EC, PO, TPP, NTPP | TPP, NTPP             | NTPP                          | DC, EC, PO, TPP, TPPi, = | DC, EC, PO, TPPi, NTPPi | TPP   |
| NTPP  | DC                      | DC                       | DC, EC, PO, TPP, NTPP | NTPP                  | NTPP                          | DC, EC, PO, TPP, NTPP    | *                       | NTPP  |
| TPPi  | DC, EC, PO, TPPi, NTPPi | EC, PO, TPPi, NTPPi      | PO, TPPi, NTPPi       | PO, TPP, TPPi, =      | PO, TPP, NTPP                 | TPPi, NTPPi              | NTPPi                   | TPPi  |
| NTPPi | DC, EC, PO, TPPi, NTPPi | PO, TPPi, NTPPi          | PO, TPPi, NTPPi       | PO, TPPi, NTPPi       | PO, TPP, NTPP, TPPi, NTPPi, = | NTPPi                    | NTPPi                   | NTPPi |
| =     | DC                      | EC                       | PO                    | TPP                   | NTPP                          | TPPi                     | NTPPi                   | =     |

Figure 2: The RCC-8 Composition Table (Cohn et al. 1997)

|       | DC      | EC      | PO      | TPP     | NTPP    | TPPi    | NTPPi   |
|-------|---------|---------|---------|---------|---------|---------|---------|
| DC    | DEPTNQt | DEPTN   | DEPTN   | DEPTN   | DEPTN   | DE      | DE      |
| EC    | DEPt    | DE      | DEPTQ   | DETN    | DE      | DE      | DE      |
| PO    | DEPt    | DEPTNQt | DEPTNQt | DEPTNQt | DEPTNQt | DEPTNQt | DEPTNQt |
| TPP   | D       | DE      | DEPTN   | TN      | N       | DEPTNQt | DEPTN   |
| NTPP  | D       | DE      | DEPTN   | TN      | N       | DEPTN   | DEPTNQt |
| TPPi  | DEPt    | DEPt    | DEPTQt  | PTQt    | PTNQt   | TQt     | TQt     |
| NTPPi | D       | DEPTN   | PQt     | TNQt    | NQt     | Qt      | Qt      |

Figure 3: The Composition Table for RCC-8 produced by ChatGPT-4. The entry in each cell uses the following coding: D (DC), E (EC), P (PO), T (TPP), N (NTPP), t (TPPi), n (NTPPi), Q (EQ). Black means that relation is correctly predicted (85 times), red means that relation is incorrectly predicted (61 times), blue means that the relation was incorrectly not predicted (49 times).

out that humans tend to have a “preferred” relation. For example, Ragni et al (2007) report on experiments performed on native German speakers and native Mongolian speakers for RCC-8. In their experiments the relations were described, but the human subjects were not allowed to draw possible configurations, so the setting is essentially equivalent to an LLM setting.

Given that humans may struggle to see all the possible relations<sup>2</sup>, determining whether there is agreement about the most preferred is good question to ask. It turns out that there is good agreement in general across and within the two cultures, with the the percentage of people agreeing with the same preferred relation ranging from 30% to 87.5% (a random choice would yield 12.5% on average since there are eight relations to choose from). (They did not query cases where the composition yields a unique relation, nor did they consider EQ as one of the two relations as this should be a trivial task.) This agreement is perhaps surprising since the two languages are linguistically very different. Ragni et al (2007) do report some differences though – for example although both language speakers preferred DC whenever it was consistent, Mongolians preferred PO over NTPPi

<sup>2</sup>The fact that some humans may struggle to compute the composition table does not stop it being a valid question to see if an LLM can determine the correct entries.

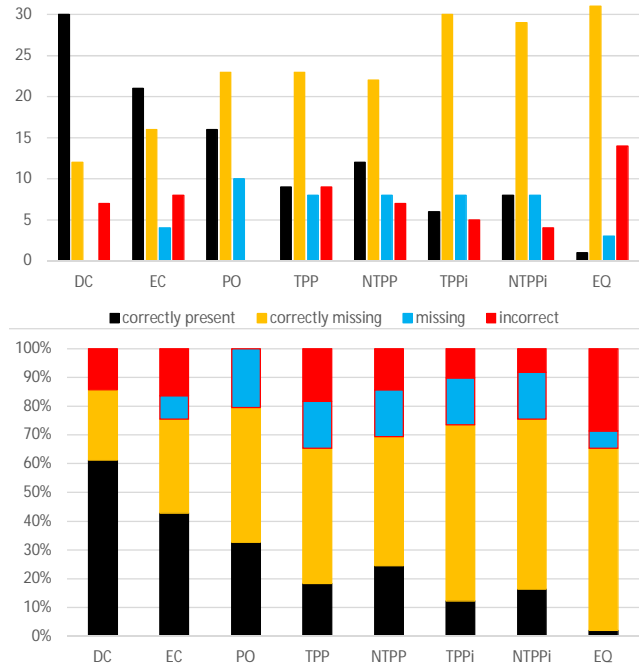


Figure 4: Relation statistics for the Composition Table for RCC-8 produced by ChatGPT-4. The upper chart shows the absolute number of relations, and the lower the relative percentage for each relation.

whereas for Germans the converse was true. Both cultures only chose EQ when composing a relation with its inverse (e.g. TPP with TPPi).

The theory of *preferred mental models* (Knauff, Rauh, and Schlieder 1995) states that people construct the simplest (computationally cheapest) model consistent with the premises. Their experiments showed that humans prefer models with the smallest overlapping complexity which explains the preference for DC noted above.

Given the difficulty reported in Experiment 1 in ChatGPT-4 correctly inferring all possible relations in a composition,

|       | DC                                 | EC                               | PO                                | TPP                                | NTPP                              | TPPi                             | NTPPi                              |
|-------|------------------------------------|----------------------------------|-----------------------------------|------------------------------------|-----------------------------------|----------------------------------|------------------------------------|
| DC    | DE P <sub>t</sub> NQ <sub>tn</sub> | DE P <sub>t</sub> N              | DE P <sub>t</sub> N               | DE P <sub>t</sub> N                | DE P <sub>t</sub> N               | DE                               | DE                                 |
| EC    | DE P <sub>tn</sub>                 | DE P <sub>t</sub> Q <sub>t</sub> | DE P <sub>t</sub> N               | DE P <sub>t</sub> N                | DE P <sub>t</sub> N               | DE P                             | DE                                 |
| PO    | D E <sub>tn</sub>                  | DE P <sub>tn</sub>               | DE P <sub>t</sub> NQ              | DE P <sub>t</sub> N                | DE P <sub>t</sub> N               | DE P <sub>t</sub> n              | DE P <sub>t</sub> n                |
| TPP   | D                                  | DE                               | DE P <sub>t</sub> N               | T <sub>n</sub>                     | N                                 | DE P <sub>t</sub> Q <sub>t</sub> | DE P <sub>t</sub> n                |
| NTPP  | D                                  | DE                               | DE P <sub>t</sub> N               | N                                  | N                                 | DE P <sub>t</sub> N              | DE P <sub>t</sub> NQ <sub>tn</sub> |
| TPPi  | DE P <sub>tn</sub>                 | DE P <sub>tn</sub>               | DE P <sub>t</sub> N <sub>tn</sub> | DE P <sub>t</sub> NQ <sub>tn</sub> | DE P <sub>t</sub> N <sub>tn</sub> | T <sub>Q</sub> <sub>tn</sub>     | DE P <sub>t</sub> N <sub>tn</sub>  |
| NTPPi | DE P <sub>tn</sub>                 | DE P <sub>tn</sub>               | DP                                | D P <sub>t</sub> N <sub>tn</sub>   | D P <sub>t</sub> N <sub>tn</sub>  | D P <sub>t</sub> N <sub>tn</sub> | D P <sub>t</sub> N <sub>tn</sub>   |

Figure 5: The Composition Table for RCC-8 produced by ChatGPT-4 using anonymised relations. The entry in each cell uses the following coding: D (DC), E (EC), P (PO), T (TPP), N (NTPP), t (TPPi), n (NTPPi), Q (EQ). Black means that relation is correctly predicted (95 times), red means that relation is incorrectly predicted (57 times), blue means that the relation was incorrectly not predicted (72 times).

asking this more specific query seems a natural task to try; this gives some measure of whether the cognitive preferences of humans correspond at all to the choices made by ChatGPT-4.

The initial prompt given to ChatGPT-4 was identical to the initial prompt above when querying about composition. In its responses, sometimes ChatGPT-4 gave a clear preference for one relation, for example:

**Prompt:** *If  $DC(x,y)$  and  $DC(y,z)$  then what is your preferred relationship between  $x$  and  $z$ ?* **Response:** *The relations  $DC(x,y)$  and  $DC(y,z)$  mean that region  $x$  is disconnected from region  $y$  and region  $y$  is disconnected from region  $z$ . These relationships don't provide any information about the spatial relationship between  $x$  and  $z$ .  $x$  and  $z$  could be disconnected, they could touch at a boundary, or they could partially overlap, among other possibilities. However, in absence of any other information and considering spatial continuity, I would say the most likely relation is  $DC(x, z)$ , assuming  $x$  and  $z$  are also disconnected. Please note that this is just a statistical assumption; it's possible for  $x$  and  $z$  to have any of the other relations defined in RCC-8.*

In other cases ChatGPT-4 justified its choice by saying it was the “cautious” choice, or the “safest choice”. The humans in the (Ragni, Tsenden, and Knauff 2007) experiment were not asked to justify their choices, but the claim made by the authors of that paper noted above is that choices were based on computational/simplicity considerations – which differs from the reasons claimed here by ChatGPT-4.

The initial prompt to ChatGPT-4 given was:

*Consider the following set of eight pairwise disjoint and mutually exhaustive binary spatial relations, between two one piece spatial 2D regions. These relations form part of the well known RCC-8 qualitative spatial reasoning calculus.  $DC(x,y)$  means that  $x$  and  $y$  are disconnected and share no spatial parts.  $EC(x,y)$  means that  $x$  and  $y$  touch at a boundary but do not share any interior parts.  $PO(x,y)$  means that  $x$  and  $y$  share a spatial part, but neither is part of the other.  $TPP(x,y)$  means that  $x$  is part of  $y$  and touches  $y$ 's*

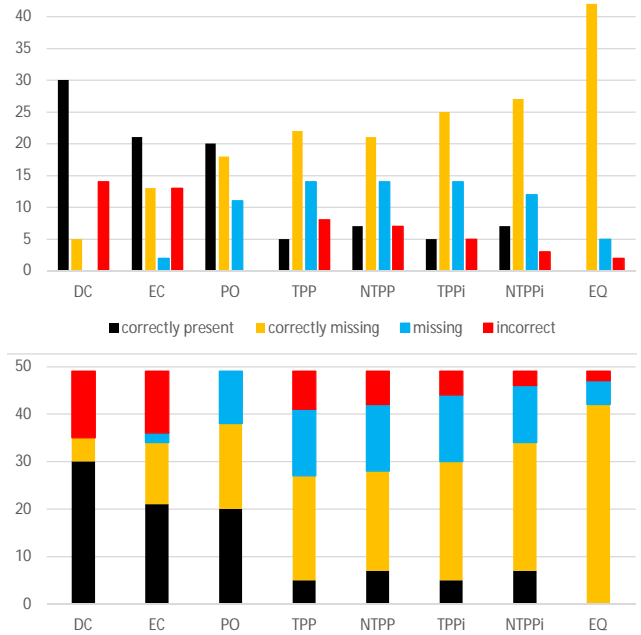


Figure 6: Relation statistics for the Composition Table for RCC-8 with anonymised relations produced by ChatGPT-4. The upper chart shows the absolute number of relations, and the lower the relative percentage for each relation.

*boundary.  $NTPP(x,y)$  means that  $x$  is part of  $y$  but does not touch  $y$ 's boundary.  $TPPi(x,y)$  is the same as  $TPP(y,x)$ .  $NTPPi(x,y)$  is the same as  $NTPP(y,x)$ . Finally,  $EQ(x,y)$  means that  $x$  and  $y$  are coincident. I will now ask you a series of questions about these relations. It is possible that in some cases there could be more than one relation that might hold; in these cases just give me your single most preferred relationship.*

As can be seen in Fig. 7, ChatGPT-4 only agreed with the average human on 20/49 (40.82%) though in a further three cases it agreed with one of the language groups (twice Mongolian, once German) giving a total of 23/49 (46.93%). In seven cases it actually predicted an impossible relation as its preferred composition<sup>3</sup>, and in all cases where the composition is unique it failed to note that. In all the other cases, 18/49 (36.63%) its preferred relation was possible but not preferred by either language group or overall.

As can be seen, sometimes ChatGPT-4 made a logically inconsistent prediction, for example with EC/NTPPi; looking at the ChatGPT-4's response in this case we can see that this is because it seems to misunderstand the meaning of NTPPi:

*If  $EC(x,y)$  and  $NTPPi(y,z)$ , this means that region  $x$  touches the boundary of region  $y$  but does not overlap with it, and region  $y$  is a part of region  $z$  but does*

<sup>3</sup>Sometimes the humans in the (Ragni, Tsenden, and Knauff 2007) experiment also predicted an impossible relation as can be seen in Table 2 of their publication.

|       | DC | EC | PO   | TPP  | NTPP | TPPi | NTPPi |
|-------|----|----|------|------|------|------|-------|
| DC    | DC | DC | DC   | DC   | DC   | DC   | DC    |
| EC    | DC | DC | EC   | EC   | EC   | EC   | EC    |
| PO    | DC | EC | PO   | PO   | PO   | PO   | PO    |
| TPP   | DC | EC | PO   | TPP  | NTPP | DC   | DC    |
| NTPP  | DC | DC | NTPP | NTPP | NTPP | DC   | DC    |
| TPPi  | DC | EC | PO   | PO   | PO   | TPP  | NTPP  |
| NTPPi | DC | DC | PO   | NTPP | NTPP | NTPP | NTPP  |

Figure 7: The table of preferred compositions for RCC-8 produced by ChatGPT-4. Green text means it agreed with the human most likely model. Orange text means it agreed with the human most likely model but failed to notify that it was the only possible model. Red means it chose an impossible relation. Blue means it chose a possible but not preferred model. A green background means that it agreed with preferred relation of one of the language groups but not the overall average.

*not touch  $z$ 's boundary, with the perspective of  $y$  being contained in  $z$ . In this scenario, it's likely that  $x$  and  $z$  are either disconnected (DC) or touch at the boundary (EC). Since  $y$ , which is part of  $z$  but does not touch  $z$ 's boundary, is touching  $x$  at the boundary, there's a possibility for indirect spatial contact between  $x$  and  $z$ . Without additional context, it's safer to assume that  $x$  and  $z$  touch at the boundary (EC). However, this is just an assumption, and the actual relationship between  $x$  and  $z$  could be different depending on the specific spatial arrangement.*

In a different case,  $TPPi/TPPi$ , the initial understanding of the relations is correct inferring that  $z$  is part of  $x$ , but then ChatGPT-4 makes a wrong step in the final part of its reasoning, and claims that  $TPP(x, y)$ :

*If  $TPPi(x, y)$  and  $TPPi(y, z)$ , this means that region  $y$  is a part of region  $x$  and touches  $x$ 's boundary, and region  $z$  is a part of region  $y$  and touches  $y$ 's boundary. Given these conditions, it's likely that  $z$  is a part of  $x$  and touches  $x$ 's boundary (TPP). This is because  $z$ , which is a part of  $y$  and touches  $y$ 's boundary, is also indirectly a part of  $x$  through  $y$ , which is a part of  $x$  and touches  $x$ 's boundary. Therefore, my preferred relationship between  $x$  and  $z$  would be  $TPP(x, z)$ , assuming that  $z$  is a part of  $x$  and touches  $x$ 's boundary. Please note that this is an assumption and the actual relationship between  $x$  and  $z$  could be different depending on the specific spatial arrangement.*

ChatGPT-4 makes exactly the same error with  $TPPi/NTPPi$ . For the wrong  $NTPPi/EC$  case, the the description of its reasoning contains a more fundamental error:

*...This is because  $y$ , which is a part of  $x$  and does not touch  $x$ 's boundary, is also only touching  $z$ 's boundary, so  $x$  and  $z$  would not have any overlap...*

A similar more fundamental error is made in the cases of  $NTPPi/TPP$ , whilst in the final two wrong cases, ChatGPT-4 misunderstands the second relation as its inverse.

In order to test whether the result was influenced by prior knowledge of RCC-8 gained as part of its training, we also

performed the same experiment, but with all the relation names prefixed by an X to disguise the connection to RCC-8. The prompt was the same as above except for the change of relation names and the omission of the second sentence. The results are given in Fig. 8. ChatGPT-4 only agreed with the average human on 24/49 (48.99%) though in a further five cases it agreed with one of the language groups (twice Mongolian, twice German), and in one case the Mongolian preference was equal between the overall choice and the choice ChatGPT-4 made giving a an agreement with the German speakers of 26/49 (53.06%) and 27/49 (55.1%) with the Mongolian speakers. In three cases (6.12%) it actually predicted an impossible relation as its preferred composition (this also happened sometimes in the human experiment), and in two cases (out of 11) where the composition is unique it failed to note that. In all the other cases, 20/49 (60.82%) its preferred relation was possible but not preferred by either language group or overall. The overall agreement with humans in this anonymised case is actually significantly higher than the non-anonymised experiment above; it is not clear if there is a reason for this and if so what.

|       | DC | EC | PO | TPP   | NTPP | TPPi  | NTPPi |
|-------|----|----|----|-------|------|-------|-------|
| DC    | DC | DC | DC | DC    | DC   | DC    | DC    |
| EC    | DC | EC | PO | EC    | EC   | EC    | EC    |
| PO    | DC | EC | PO | PO    | PO   | PO    | PO    |
| TPP   | DC | EC | PO | TPP   | NTPP | DC    | DC    |
| NTPP  | DC | DC | PO | NTPP  | NTPP | DC    | DC    |
| TPPi  | DC | EC | PO | PO    | PO   | TPPi  | NTPPi |
| NTPPi | DC | DC | PO | NTPPi | NTPP | NTPPi | NTPPi |

Figure 8: The table of preferred compositions for RCC-8 produced by ChatGPT-4 with anonymised relations. The relation names are coloured in the same way as in Table 7.

### Experiment 3: Spatial Continuity

Continuity networks were introduced by Randell and Cohn(1989) to represent the set of possible 'next' relations that might obtain, assuming that motion is continuous and any transformations in the shape and/or size of an object are also continuous. Subsequently these have been termed "conceptual neighbourhoods" (Freksa 1992). Fortunately, the abbreviation for both of these terms is CN, which we shall use henceforth. The CN for RCC-8 is depicted in Fig. 1. The initial prompt was the same as for the composition table prompt above, but with two additional sentences at the end: *In all these questions, assume that all motion is continuous, and any deformations of the spatial regions are also continuous. Also please describe the motion or changes of size or shape that would be required for the change of relation to take place.* Each subsequent prompt also emphasises the need for the relation to be an immediate next relation. As an example prompt/response pair consider the EC case:

**Prompt:** *If  $EC(a, b)$  holds, then what are the next possible relations that might hold as the two regions translate or change their shape or size? I am interested in the immediate next relation, so no intermediate relation must hold.* **Response:** *If  $EC(a, b)$  holds,*



which means that the two regions  $a$  and  $b$  are just touching at a boundary but do not share any interior parts, then the immediate next possible relations as the two regions translate or change their shape or size could be either  $DC(a,b)$  or  $PO(a,b)$ .  $DC(a,b)$ : This would occur if the regions move away from each other and are no longer in contact at all, becoming fully disconnected.  $PO(a,b)$ : This would occur if one region begins to overlap with the other, but neither region is completely contained within the other. They would then be partially overlapping, sharing some but not all of their spatial parts. The transitions to these relations would involve either moving the regions apart or moving them further into each other. In the case of transition to  $PO(a,b)$ , there could also be some deformation of the regions if they change shape to allow for the partial overlap.

An example of a faulty prediction is from  $PO$ :

If  $PO(a,b)$  holds, indicating that regions  $a$  and  $b$  partially overlap but neither is fully contained within the other, the immediate next relations could be  $EC(a,b)$ ,  $DC(a,b)$ ,  $TPP(a,b)$ ,  $TPPi(a,b)$ , or  $EQ(a,b)$ .  $EC(a,b)$ : This would occur if the overlapping parts of  $a$  and  $b$  are reduced to just boundary parts, meaning the regions are touching but no longer overlapping.  $DC(a,b)$ : This would occur if regions  $a$  and  $b$  are moved or deformed such that they no longer overlap or touch at all.  $TPP(a,b)$  or  $TPPi(a,b)$ : This would occur if one region becomes fully contained within the other, either by moving into it or shrinking.  $TPP(a,b)$  means  $a$  is part of  $b$  and touches  $b$ 's boundary.  $TPPi(a,b)$  means  $b$  is part of  $a$  and touches  $a$ 's boundary.  $EQ(a,b)$ : This could occur if the regions are reshaped or resized such that they coincide exactly, becoming the same region in space. The transitions to these relations would involve either moving the regions, resizing them, or reshaping them in some way.

Here, ChatGPT-4 has mistakenly said that  $DC$  is a next relation, failing to notice that it has to pass through  $EC$  on the way.

As can be seen, the predictions are mostly correct with just five errors: three where a link is falsely predicted and two where a link is missing. The two missing links are both from  $EQ$  (to  $NTPP$  and to  $NTPPi$ ). Interestingly the links in the reverse direction are correctly predicted, so ChatGPT-4 is not able to reason that a link in one direction implies the link in the other direction should also be present. Similarly whilst the absence of a link from  $DC$  to  $PO$  is correctly predicted, the inverse case is not, and is also the case for  $PO$  to/from  $NTPP$  and  $NTPPi$ .

In order to test whether the result was influenced by prior knowledge of  $RCC-8$  gained as part of its training, we also performed the same experiment, but with all the relation names prefixed by an  $X$  to disguise the connection to  $RCC-8$ . The prompt was the same as above except for the change of relation names and the omission of the second sentence. The results are given in Fig. 10. There are 3 incorrectly predicted

|       | DC | EC | PO | TPP | NTPP | TPPi | NTPPi | EQ |
|-------|----|----|----|-----|------|------|-------|----|
| DC    |    | x  |    |     |      |      |       |    |
| EC    | x  |    | x  |     |      |      |       |    |
| PO    | x  | x  |    | x   |      | x    |       | x  |
| TPP   |    |    | x  |     | x    |      |       | x  |
| NTPP  |    |    | x  | x   |      |      |       | x  |
| TPPi  |    |    | x  |     |      |      | x     | x  |
| NTPPi |    |    | x  |     |      | x    |       | x  |
| EQ    |    |    | x  | x   |      | x    |       |    |

Figure 9: The Continuity Table for  $RCC8$  produced by ChatGPT-4. An ‘x’ means that the relation in that column is predicted as an immediate neighbour of the relation in that row. An empty box means that the relation is not predicted as an immediate neighbour. Green means that the prediction was correct and red that it was incorrect. The leading diagonal is white since a relation is not a next relation of itself.

links, 3 missing links, 19 correctly predicted links and 31 correct missing links, giving an accuracy of 50/56 (89.2%). This is slightly worse than the case above. There is one more missing link but the missing links are all different in the two cases. Although there are the same number of wrong links, only one of these is in common ( $PO$  to  $DC$ ). Overall the results are broadly similar and may be due to the stochastic nature of ChatGPT-4’s responses, suggesting that either the disguise was not very effective, or that prior training did not really affect the response and it was able to reason from ‘first principles’ (if not always correctly) in response to each prompt.

|       | DC | EC | PO | TPP | NTPP | TPPi | NTPPi | EQ |
|-------|----|----|----|-----|------|------|-------|----|
| DC    |    | x  |    |     |      |      |       |    |
| EC    | x  |    | x  |     |      |      |       |    |
| PO    | x  | x  |    | x   |      | x    |       |    |
| TPP   |    | x  | x  |     | x    |      |       |    |
| NTPP  |    |    |    | x   |      |      |       | x  |
| TPPi  |    | x  | x  |     |      |      | x     |    |
| NTPPi |    |    |    |     |      | x    |       | x  |
| EQ    |    |    | x  | x   | x    | x    | x     |    |

Figure 10: The Continuity Table for  $RCC-8$  produced by ChatGPT-4 using disguised relation names. The meaning of the colouring is the same as in Fig. 9.

## Concluding Remarks and Future Work

This investigation has supported the widely-held view that LLMs can struggle to do reasoning tasks<sup>4</sup>. In the case of Experiment 1, in which ChatGPT-4 was asked to compute the entire composition table for  $RCC-8$ , this is a non trivial task even for humans, so it is perhaps not surprising that ChatGPT-4 did not achieve 100% accuracy – the scores of 71.94% (and 67.09% for the anonymised relations) are clearly much better than chance and do suggest a reasonable facility to perform such computations. A detailed analysis of the actual conversations in the supplementary material shows that

<sup>4</sup>Bender et al (2021) have observed that LLMs might be regarded just as “stochastic parrots” and thus it is not surprising that precise, logically correct deductive reasoning is challenging for an LLM.

sometimes ChatGPT-4 does appear able to do some interesting (qualitative) spatial reasoning, but often fails, sometimes making elementary mistakes. It also shows inconsistency in being able to reason correctly about a relation but not its inverse. It also sometimes confuses a relation with its inverse. It is possible that fine tuning, explicit chain-of-thought prompting, or more carefully engineered prompts might improve performance; however, given the stochastic nature of LLMs it seems unlikely that the results would be as good as logical reasoning (the experiment on preferred relations is of course not strictly a logical reasoning exercise, except for the requirement not to predict spatially impossible relations).

There are a variety of avenues for further work which present themselves. Other calculi could be experimented with – for example the coarser calculus RCC-5, or calculi for reasoning about direction or size (Cohn and Renz 2008). Other LLMs could be evaluated – though since new LLMs and new LLM versions are continually being released, this is a challenge with no definite stopping point. Tracking the change in performance of a particular LLM across releases would also be of interest – though in the case of closed LLMs such as ChatGPT-4 where the owners have the right to harvest user conversations and use them for future training, it will not be clear if any improvement is the result of leakage from the previous conversation or more general performance improvement<sup>5</sup>. It has already been observed (Cohn and Hernandez-Orallo 2023) that different LLMs have different strengths – determining which LLMs are better at which spatial reasoning tasks would also be worth of future investigation. The overall conclusion that LLMs in general struggle with more complex spatial reasoning tasks is likely to remain the case, at least for the foreseeable future. In the API version of GPT, different temperatures could be tried, and multiple runs with averages computed. Different prompts and prompting strategies could be tried, though arguably since QSR has always been viewed as a form of commonsense reasoning, it should not be necessary to devise specific prompts to elicit commonsense behaviour.

It is not clear how successful the anonymisation was – in one case I mistyped an X relation and it was able to suggest the intended relation name, suggesting that it has the ability to dissect relation names; thus more sophisticated anonymisation might be tried. In earlier work (Cohn and Hernandez-Orallo 2023) we had already done some limited experimentation asking an LLM to reason about spatial relations in a real world context rather than the purely abstract setting used in the experiments in this paper – it would be interesting to conduct more extensive tests LLMs doing compositional reasoning in a more realistic setting, and similarly for the continuity experiment.

Experiment 2 above already investigated how LLM performance compared to human performance to a limited extent but further investigation would be worthwhile, including a head-to-head comparison rather than simply taking a result from the literature originally intended to investigate a different question. Another interesting avenue for further

---

<sup>5</sup>However, note that no feedback was given to ChatGPT-4 as to whether the proffered response was correct or not.

work will be to explore the use of multimodal FMs – when humans perform spatial reasoning tasks including the challenge of building a composition table, it is natural to use pencil and paper to sketch diagrams and possible scenarios – investigating whether a multi-modal FM with such abilities (including the ability to analyse its own drawings) would be of great interest to the spatial reasoning community.

As mentioned above, another possible avenue of research is to investigate different prompting strategies, including k-shot (Dang et al. 2022), chain-of-thought(Wei et al. 2022) and tree-of-thought(Yao et al. 2023) strategies. Not doing so was deliberate in this paper as I was interested in exploring in how the “vanilla” LLM would perform. Whilst for specific downstream tasks, fine-tuning or employing specific prompting strategies may be reasonable, there is an argument to be made that for commonsense reasoning, this is not a reasonable strategy since the task is a general one rather than a specific downstream task.

## Data statement

All the conversations with ChatGPT-4 that support the summary tables in this paper can be found at <http://tinyurl.com/qz23sup>.

## Acknowledgments

This work was supported by: The Alan Turing Institute; the Economic and Social Research Council (ESRC) under grant ES/W003473/1; the Turing’s Defence and Security programme through a partnership with the UK government in accordance with the framework agreement between GCHQ and The Alan Turing Institute.

## References

- Bender, E. M.; Gebru, T.; McMillan-Major, A.; and Shmitchell, S. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 610–623.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33: 1877–1901.
- Cai, Z.; Chang, B.; and Han, W. 2023. Human-in-the-Loop through Chain-of-Thought. *arXiv preprint arXiv:2306.07932*.
- Chen, J.; Cohn, A. G.; Liu, D.; Wang, S.; Ouyang, J.; and Yu, Q. 2015. A survey of qualitative spatial representations. *The Knowledge Engineering Review*, 30(1): 106–136.
- Cohn, A. G.; Bennett, B.; Gooday, J.; and Gotts, N. M. 1997. Qualitative Spatial Representation and Reasoning with the Region Connection Calculus. *Geoinformatica*, 1(3): 275–316.
- Cohn, A. G.; and Hazarika, S. M. 2001. Qualitative spatial representation and reasoning: An overview. *Fundamenta Informaticae*, 46(1-2): 1–29.

- Cohn, A. G.; and Hernandez-Orallo, J. 2023. Dialectical language model evaluation: An initial appraisal of the common-sense spatial reasoning abilities of LLMs. *arXiv preprint arXiv:2304.11164*.
- Cohn, A. G.; and Renz, J. 2008. *Handbook of Knowledge Representation*, chapter Qualitative Spatial Representation and Reasoning. Elsevier.
- Dang, H.; Mecke, L.; Lehmann, F.; Goller, S.; and Buschek, D. 2022. How to Prompt? Opportunities and Challenges of Zero- and Few-Shot Learning for Human-AI Interaction in Creative Applications of Generative Models. *arXiv:2209.01390*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Freksa, C. 1992. Temporal reasoning based on semi-intervals. *Artificial intelligence*, 54(1-2): 199–227.
- Gerevini, A.; and Renz, J. 2002. Combining topological and size information for spatial reasoning. *Artificial Intelligence*, 137(1-2): 1–42.
- Guesgen, H. W. 1989. Spatial reasoning based on Allen’s temporal logic, Technical Report TR-89-049. Technical report, International Computer Science Institute Berkeley.
- Jonsson, P.; and Drakengren, T. 1997. A Complete Classification of Tractability in RCC-5. *J. Artif. Intell. Res.*, 6: 211–221.
- Knauff, M.; Rauh, R.; and Schlieder, C. 1995. Preferred mental models in qualitative spatial reasoning: A cognitive assessment of Allen’s calculus. In *Proceedings of the seventeenth annual conference of the Cognitive Science Society*, 200–205.
- Moratz, R. 2006. Representing relative direction as a binary relation of oriented points. In *ECAI*, volume 6, 407–411.
- Mukerjee, A.; and Joe, G. 1990. A Qualitative Model for Space. In Shrobe, H. E.; Dietterich, T. G.; and Swartout, W. R., eds., *Proceedings of the 8th National Conference on Artificial Intelligence. Boston, Massachusetts, USA, July 29 - August 3, 1990, 2 Volumes*, 721–727. AAAI Press / The MIT Press.
- Ragni, M.; Tseden, B.; and Knauff, M. 2007. Cross-cultural similarities in topological reasoning. In *Spatial Information Theory: 8th International Conference, COSIT 2007*, 32–46. Springer.
- Randell, D.; and Cohn, A. 1989. Modelling topological and metrical properties in physical processes. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, 357–368.
- Randell, D.; Cui, Z.; and Cohn, A. G. 1992. A Spatial Logic Based on Regions and Connection. In *3rd International Conference on Knowledge Representation and Reasoning, 1992*, 165–176.
- Roumeliotis, K. I.; and Tselikas, N. D. 2023. ChatGPT and Open-AI Models: A Preliminary Review. *Future Internet*, 15(6): 192.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Chi, E.; Le, Q.; and Zhou, D. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.
- Yao, S.; Yu, D.; Zhao, J.; Shafraan, I.; Griffiths, T. L.; Cao, Y.; and Narasimhan, K. 2023. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*.