



Physical models and embodied cognition

Ulrich E. Stegmann¹

Received: 23 March 2018 / Accepted: 31 August 2018 / Published online: 10 September 2018
© The Author(s) 2018

Abstract

Philosophers have recently paid more attention to the physical aspects of scientific models. The attention is motivated by the prospect that a model's physical features strongly affect its use and that this suggests re-thinking modelling in terms of extended or distributed cognition. This paper investigates two ways in which physical features of scientific models affect their use and it asks whether modelling is an instance of extended cognition. I approach these topics with a historical case study, in which scientists kept records not only of their findings, but also of some of the mental operations that generated the findings. The case study shows how scientists can employ a physical model (in this case diagrams on paper) as an external information store, which allows alternating between mental manipulations, recording the outcome externally, and then feeding the outcome back into subsequent mental manipulations. The case study also demonstrates that a model's physical nature allows replacing explicit reasoning with visuospatial manipulations. I argue, furthermore, that physical modelling does not need to exemplify a strong kind of extended cognition, the sort for which external features are mereological parts of cognition. It can exemplify a weaker kind, instead.

Keywords Task decomposition · Visuospatial reasoning · Mental rotation · Crick · Gamow · Protein synthesis

Recent work on scientific models has emphasised their material nature (e.g. Giere 2002; Knuuttila 2011; Kuorikoski and Ylikoski 2015). Even abstract models are said to be physical insofar as the notation systems used to explore them involve physical marks, e.g. mathematical equations on a piece of paper (Knuuttila 2011). The emphasis on a model's physicality is motivated by the view that the efficiency of scientific modelling hinges on the extent to which models have a physical reality, as opposed to being merely imagined. Since the use of models is an important part of understanding their role in science, it is important to understand how a model's physicality affects its

✉ Ulrich E. Stegmann
u.stegmann@abdn.ac.uk

¹ School of Divinity, History and Philosophy, University of Aberdeen, 50-52 College Bounds, High Street, Aberdeen AB24 3UB, UK

use.¹ In addition, it has been argued that the physicality of models requires re-thinking the cognitive foundations of modelling practice. Arguably, modelling is an instance of “distributed cognition” (Giere 2002; Nersessian 2006).

This paper addresses two questions. First, how exactly does the physicality of models affect their use? According to one answer, physicality enables models to be employed as external information stores. On this view, models are “material objects that can be used to keep score of inferential moves and in this way be used to help in reasoning about the phenomena of interest” (Kuorikoski and Ylikoski 2015, p. 3818). Although plausible in principle, demonstrating that models can actually be employed for this purpose requires showing, in at least one particular example, not only that scientists keep external records of their reasoning, but also that they retrieve and consult their records later and that the stored information subsequently affects their reasoning. This paper attempts such a demonstration. Furthermore, we must ask whether acting as an information store is the only contribution of physicality or whether there are others. I will argue that there is at least one additional benefit, namely replacing explicit reasoning with visuospatial manipulations.

The second question concerns the relation between scientific modelling and distributed cognition. Emphasising the physicality of models often goes along with construing modelling as an instance of “distributed” or “extended cognition” (Giere 2002; Nersessian 2006; Kuorikoski and Ylikoski 2015). On a strong reading, distributed cognition is the view that the external resources on which our cognitive processes rely are as much part of these processes as the relevant neural events inside our brains (e.g. Menary 2010; Wilson and Foglia 2015). Accordingly, when scientists employ physical models, their cognitive processes literally extend into their surroundings, encompassing the physical models and scientists’ interactions with them. But just how tight is the link between distributed cognition in this strong sense and scientific models? I will argue that it is not tight at all, not even when modelling relies heavily on manipulating physical marks. Instead, a weaker and less controversial form of distributed cognition may well be sufficient for understanding modelling with physical aids.

The paper is organised as follows. Section 1 introduces the case study. The case study is a diagrammatic method used in the 1950s for refining a hypothesis about protein synthesis. I first introduce the hypothesis and the associated visual representations (diamond diagrams). I then show how these diagrams were used as physical models in order to elaborate and, eventually, refute the hypothesis. Section 2 addresses the effects of a model’s physicality. I first consider the use of physical models as information stores and then as supporting visuospatial reasoning. Section 3 addresses extended cognition.

1 The case study

1.1 Gamow’s DNA cavities and their diagrammatic representation

Gamow (1954a) proposed a mechanism for the linear ordering of individual amino acids in proteins. For each of the 20 amino acids, Gamow postulated a unique DNA

¹ It is also plausible that different physical features affect modelling in different ways. For example, it may be important whether the physical features are pencil marks on paper or visualisations on a computer screen. However, I do not pursue this issue here.

cavity. The cavities were arranged in a linear chain. They would attract their specific amino acids by means of a lock-and-key mechanism and hold them in place until they could be joined up into a chain. The chain would then detach from the DNA. This is the hypothesis that was subsequently investigated by diagrammatic means (“Gamow’s hypothesis” or “proposal” in what follows). Although Gamow postulated 20 distinct cavities for the 20 amino acids, he left open which cavity specifies which amino acid.

Let us look more closely at the cavities’ structure and their diagrammatic representations. Cavities were thought to be composed of four bases arranged in a cross pattern, thereby creating diamond-shaped holes on the surface of the double helix. The cavities were depicted as diamonds with numerals (bases) in their corners (Fig. 1). A single cavity was depicted as a single diamond, and a sequence of cavities as a sequence of diamonds (Fig. 2). Note that Fig. 2 lacks the straight lines of the diamonds depicted in Fig. 1. However, the four numerals belonging to one diamond are still arranged in the shape of a diamond, just as in Fig. 1.

1.2 Employing diagrams for elaborating and testing a hypothesis

In the first instance, diamond diagrams were static depictions of purported real-world entities (DNA cavities). In addition, diamond diagrams were malleable devices for filling gaps in Gamow’s (1954a) original proposal. One scientist engaged in this activity was Francis Crick. His unpublished notes have survived and now allow us to understand the diagrammatic procedure (Crick 1953–1954). I will explain some of Crick’s work by reference to a handwritten note that remained unpublished (pp. 17–26 in PDF).² The note counts 9 numbered pages, but we only need the brief section reprinted in Fig. 2. This section contains three main steps. First, Crick wrote down a known amino acid sequence. He then wrote down the diamonds he had assigned to specific amino acids in previous sections of his note. The third step was key: it was the attempt to derive new assignments from the information contained in the diagrams. Let us look at these steps in more detail.

For the first step, Crick took the sequence of insulin B (Sanger and Tuppy 1951) and represented its middle section as a series of three-letter words (e.g. *Leu* for leucine, Fig. 2). On Gamow’s hypothesis, the existence of this particular amino acid sequence implied the existence of a corresponding series of DNA cavities, which could synthesize it. If Gamow was right, then it should be possible to represent the insulin sequence as a sequence of the corresponding DNA cavities and, hence, as a series of diamonds.

However, translating the insulin sequence into a diamond chain was far from trivial. Gamow had proposed that adjacent cavities share bases such that a given cavity shares two bases with its upper and two with its lower neighbour. Since bases were represented by numerals, only diamonds that had the same two numerals on a side could be written next to one another (Fig. 2, circled numerals). Crick approached the task by first focusing on the diamonds he had already allocated to amino acids. In earlier sections

² This is probably the document to which Crick referred later as “proof”: “I showed some time ago that the B chain could not be coded, but the proof is long and intricate and not worth reproducing” (Crick 1955, p. 2). The nine pages are indeed “long and intricate”, and Crick provided several other and simpler objections against Gamow’s proposal.




Fig. 1 Table of 20 diamonds. Each diamond depicted a DNA cavity that (purportedly) specified one of the 20 amino acids. Numerals denote DNA bases; the horizontal lines within diamonds connect complementary bases; capital letters denote diamonds (Gamow 1954b). Reproduced with the permission of The Royal Danish Academy of Sciences and Letters

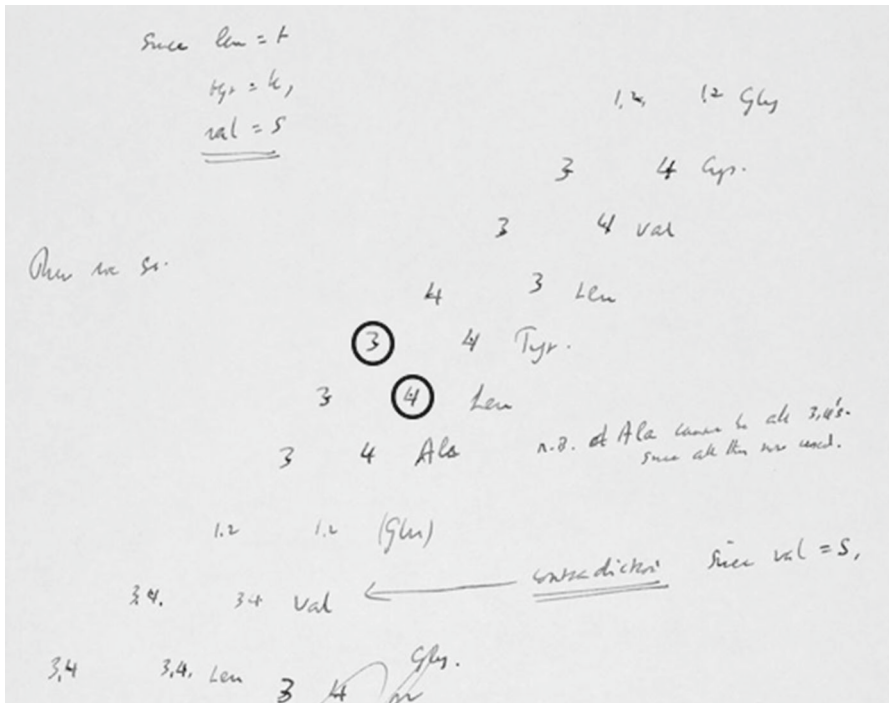


Fig. 2 Modelling DNA cavities by means of manipulating diamonds. A section of the amino acid sequence of insulin B is represented as (1) a series of three-letter words, e.g. *Leu*, *Tyr*, ... and as (2) a series of diamonds. The encircled numerals depict the bases shared by the leucine- and tyrosine-specifying cavities. Extract from p. 9 of Crick’s notes on “Gamow’s code” (Crick 1953–1954, p. 26 of PDF). Published with the permission of the Wellcome Collection

of the 9-page note, he had allocated the *T*-, *K*-, and *S*-diamonds of Gamow’s (1954b) table³ to leucine, tyrosine, and valine, respectively (Fig. 1).⁴ The assignments were noted at the top of Fig. 2 (“*leu*=*t*, *tyr*=*k*, *val*=*s*”). Insulin B contained a short sequence entirely composed of these amino acids, i.e. the sequence leucine-tyrosine-leucine-valine. Since all these amino acids had been allocated to diamonds, their sequence could be represented as a sequence of diamonds (*T-K-T-S*).

³ Gamow published two tables, the original one with letters J and Q (Gamow 1954a) and a later one without those letters (Gamow 1954b). Crick used the latter.

⁴ The tyrosine- and leucine-specifying cavities could only be represented by diamonds that could be placed next to one another by sharing their numerals. Examples for such diamonds are *K* and *T*. The *K*-diamond’s lower left side is defined by numerals 3 and 4, as is the upper right side of the *T*-diamond (Fig. 1). One can therefore write a *K*-diamond above a *T*-diamond. It followed that the *K*- and *T*-diamonds might represent the tyrosine- and leucine-specifying cavities, respectively. By way of contrast, consider the *M*-diamond. The *M*-diamond’s lower left side is defined by numerals 1 and 1 (Fig. 1), and so an *M*-diamond cannot be written above *T*. Consequently, if the *T*-diamond depicts the leucine-specifying cavity, then the *M*-diamond cannot depict the tyrosine-specifying cavity (and vice versa). In this way, the amino acid sequence of insulin B helped identify the assignments between cavities and amino acids: the *K*-diamond might specify the amino acid tyrosine, but the *M*-diamond could not.

The third step aimed at establishing *new* diamond-amino acid assignments. This step can be subdivided further. The start was made with an implicit rule according to which the numerals in the left–right corners of a given diamond could only be 3–4, 4–3, 1–2, or 2–1. The rule reflected the convention that the left–right numerals depicted the complementary bases of the corresponding DNA cavity. Consider the diamond marked ‘Cys’ in Fig. 2 and note the 3 in its left corner. For a 3 on the left, the rule prescribed a 4 on the right. Hence a 4 could be added to the right corner of that diamond. Similarly, the diamond marked ‘Ala’ required a 3 in its left corner since it had a 4 on the right (Fig. 2). After labelling the horizontal corners, attention shifted to the bottom corner of the diamond marked ‘Ala’.⁵ Its bottom corner could only be 1 or 2 (“1, 2” Fig. 2).⁶ Given the numerals in its other corners, that diamond could be either *D* or *N*. Since the diamond needed to depict an alanine-specifying cavity (due its position within the chain), that type of cavity could only be represented by either *D* or *N*.

While translating the insulin B sequence one-by-one into the corresponding DNA cavities, Crick eventually reached the valine in the sequence valine-glutamic acid-alanine-leucine-tyrosine-leucine. The only diamond left to represent the valine-specifying cavity was *S*. Yet drawing an *S*-diamond in that position proved impossible. An *S*-diamond would need a 3 in its top corner (Fig. 1), but the top corner of the diamond in that position of the chain could only be 1 or 2.⁷ This result allowed a conclusion about the corresponding cavity. Since one could not draw an *S* in the right position of the diamond chain, there could not be a valine-specifying cavity where there needed to be one. However, the empirical data (i.e. the presence of a valine) clearly implied that on Gamow’s proposal there had to be such a cavity. Crick highlighted this result with an arrow, adding “contradiction since val = s” (Fig. 2). The impossibility of drawing an *S*-diamond was therefore an important result. It meant that insulin B could not be synthesized in the way Gamow had proposed. His proposal was unworkable.

1.3 Mechanistic modelling with diagrams

Before turning to the physicality of models, let us pause to consider why the diagrammatic procedure was an instance of mechanistic research and modelling.

⁵ Attention also shifted to the top corner of the diamond marked ‘Cys’, but I will not trace it further.

⁶ Numerals 3 and 4 were excluded for the following reasons. If one labels the bottom corner 3 and then compares the resulting diamond with the other diamonds in the chain, it turns out that the diamond would be an *S*. If one assigns 4, then the diamond would be *T*. Now, *S* and *T* had already been assigned to valine and leucine, respectively, and therefore were unavailable to represent any other amino acids. The only numerals generating unassigned diamonds were 1 (yielding a *D*) and 2 (yielding an *N*).

⁷ The top corner (of the diamond that was supposed to represent the valine-specifying cavity) could only be 1 or 2 for the following reasons. (1) The top corner is identical with the left corner of the adjacent diamond, which represented the glutamic acid-specifying cavity, and that corner was 1 or 2. (2) The corner was 1 or 2, because the right corner of that diamond was 2 or 1 (1–2 complementarity rule). (3) The right corner was 2 or 1, because it could be neither 3 nor 4. (4) It could not be 3 because if it was 3, then the adjacent diamond would be an *S*-diamond, which Crick had stipulated as representing the valine-specifying cavity (but the adjacent diamond needed to represent an alanine-specifying cavity). (4) It could not be 4 because if it was 4, then the adjacent diamond would be a *T*-diamond, which Crick had stipulated as representing the leucine-specifying cavity (rather than the required alanine-specifying cavity).

The diagrammatic method was an instance of *modelling* because of the way diagrams were modified and then enabled inferences to purported real-world entities. Consider the diamond that depicted the alanine-specifying cavity (Fig. 2). Initially the diamond lacked numerals in its left and bottom corners. It was then subjected to a series of operations, such as considering the already labelled corners and the other diamonds in the chain. On this basis, the missing numerals were added and it was concluded that the diamond in this position was either *D* or *N*. Since the diamond in this position had to depict the alanine-specifying cavity, it was possible to draw the following inference about a real-world target: the alanine-specifying cavity has one of two possible base configurations, either the one depicted by the *D*- or the one depicted by the *N*-diamond. In short, the diagrammatic method proceeded by modifying certain initial features of one system (diagrams) and then taking the resulting features as the basis for justifying conclusions about a real-world system (template units). And this procedure is the strategy of working with scientific models (Frigg and Hartmann 2012). The more general point here is that, while Gamow's article presented his mechanistic hypothesis in both linguistic and visual formats (text and diagrams, respectively), his modelling practice employed the visual format.

Let us turn to the mechanistic character of the diagrammatic method. Employing the method amounted to *mechanistic* research because it was supposed to help elucidate certain components of a mechanism as well as elaborate the corresponding mechanistic hypothesis. More generally, the aim was to determine the mechanism responsible for arranging amino acids into a linear order (the explanandum phenomenon). Gamow (1954a) hypothesized that the linear order of amino acids was the result of incoming amino acids being attached to a pre-existing template, i.e. to the linear series of cavities on the surface of the DNA double helix. Thus, Gamow postulated a number of entities and activities, e.g. cavities and non-covalent bonding, that could give rise to the explanandum phenomenon. Since Gamow's hypothesis concerned a possible mechanism for arranging amino acids, it was a mechanistic "how-possibly model" in Craver's (2008) sense.

The mechanistic character of the diagrammatic method is also apparent from the specific manner in which it helped to elaborate Gamow's hypothesis, a manner seen in much mechanistic research. Mechanistic investigations often seek to identify the physical make-up of functionally characterised components (e.g. Craver and Darden 2013). Scientists may know or suspect how a (hypothetical) component functions, e.g. what it contributes to the operation of a mechanism, while at the same time being in the dark about the component's physical identity; subsequent research then aims at determining the identity. The diagrammatic method was used for exactly this purpose. Gamow (1954a) stipulated the basic physical structure of DNA cavities as well as their function (i.e. to specify certain amino acids). But he had no idea about the specific structure of, say, the alanine-specifying cavity. Determining this structure was the goal of the diagrammatic method. Now, the method fell somewhat short insofar as it failed to assign *one* of the 20 configurations (diamonds) to, say, the alanine-specifying cavity. Nevertheless, it did identify *two* possible configurations for several functionally identified cavities, e.g. for the alanine-specifying cavity and the upper neighbour of *S*. For these cavities, the method allowed reducing the number of possible configurations from 20 to 2. The diagrammatic method therefore elaborated and refined Gamow's

original hypothesis by making progress in determining the specific, physical structure of functionally individuated components.⁸

The ‘dry’ and non-experimental character of manipulating diagrams may appear unusual for a mechanistic method, since mechanistic research is normally associated with extensive experimental and laboratory work. But the diagrammatic method of the 1950s was no less aimed at elucidating the mechanism of protein synthesis than the biochemical procedures of the 1960s.

2 How did the diagrams’ physicality affect their use?

2.1 Task decomposition and external information stores

Reasoning with external diagrams is known to have several benefits, such as freeing up working memory by storing information externally (Tversky 2005; Hegarty and Stull 2012). External diagrams are especially useful to the extent they allow decomposing a complex task into a series of smaller, more manageable activities. In task decomposition, one engages mental imagery in a stepwise fashion by recording intermediate results externally before proceeding to the next mental task (Hegarty et al. 2015). Crick employed the diagrammatic method in exactly this manner. The task of assigning amino acid-cavities proceeded in three major steps, each of which involved mental processing followed by recording intermediate results externally: (1) identify a known amino acid sequence and write it down, (2) identify the DNA cavities that are already assigned to diamonds and then writing out the series of diamond, and (3) attempting new assignments. The latter task was, in turn, decomposed into steps in which cognitive processing alternated with recording intermediate results.

The first step in attempting new assignments involved recalling the rule that horizontal corners could only be labelled in pairs of 1–2 or 3–4, applying the rule to a given diamond, and adding the appropriate numeral to the corners. The second step was to mentally add one of the numerals to the bottom corner and assess whether the resulting diamond would be identical to any of the fully labelled diamonds in the chain. If the diamonds were distinct, then the numeral was added to the corner (if they were identical, the process was repeated with a different numeral). Given the psychological benefits of employing external diagrams, it is not difficult to conjecture why the diagrammatic method was useful. In the last step, for instance, a derived diamond had to be compared with the previously labelled diamonds in the chain. The fact that previous diamonds had been recorded on paper made this comparison quicker and more reliable than if one relied on memory. For similar reasons it was beneficial to write down the numerals for the horizontal corners. For then one did not need to remember them when comparing the diamond to others. In conclusion, the diagrammatic method

⁸ The numerals of a diamond constrained not only the numerals of its direct neighbours, but also those of its near-next neighbours. For instance, the *T*-diamond implied that its direct neighbour be *D* or *N*. These diamonds had a 1 and 2 in their bottom corners, respectively. The lower neighbour of a *D*- or *N*-diamond therefore required either 1 or 2 in its top corner. This fact eliminated 6 of the 20 possible diamonds as the lower neighbours of *D* and *N*. Consequently, the number of cavities that could specify glutamic acid was reduced from 20 to 14.

allowed task decomposition, including the conversion of mental images to external marks for later reference. The diagrams thus served as an external information store and freed up working memory. As noted in the introduction, what matters here is that the diagrams were physical as opposed to merely imagined. Not at issue here is the specific type of physical medium used, i.e. ink and paper instead of, say, chalk and blackboard.⁹

The phenomenon of task decomposition sheds light on the relative significance of physical diagrams and their mental images. If diagrammatic reasoning involves both, then what is their relation? Is one somehow more important than the other? Or do they play distinct roles and, if so, which? These questions form the backdrop of a dispute about diagrammatic reasoning with chemical formulas. Whereas Goodwin (2008) argued that work with chemical formulas relies on manipulating formulas on paper, Perini (2012) objected that chemists mostly manipulate mental images of formulas. The present case study suggests a possible resolution. In the case of diamond diagrams, researchers manipulated a mental image of a diagram, recorded the result on paper, and then iterated this process until all viable options were exhausted. Thus, the final diagram materialised gradually as a result and record of mental manipulations. Both physical diagrams and their mental images were important, but their roles were distinct: physical diagrams acted as an external information store, whereas their mental images were the substrate of mental manipulations. This may also be the pattern at work in the case of chemical formulas.

2.2 Visuospatial reasoning

2.2.1 Functional equivalence

Reasoning with external diagrams can have a second major benefit, which is to replace explicit reasoning by something less cognitively demanding, i.e. visuospatial processing (Tversky 2005; Hegarty and Stull 2012). This benefit, too, is manifest in the diagrammatic method. At this point, another feature of Gamow's proposal becomes relevant.

Gamow (1954a) assumed that some cavities were functionally equivalent while being structurally distinct: some cavities could specify the same kind of amino acid despite having different base configurations. In particular, cavities differing only in the position of their complementary base pair specified the same amino acid. The pair 3–4, for instance, could be positioned in two ways. In one cavity, base 3 might sit on the DNA parent strand and 4 on the daughter strand (top diamond in Fig. 3). The reverse might hold in another cavity (second diamond in Fig. 3). Gamow (1954b) argued that this structural difference was functionally irrelevant because most amino acids were bilaterally symmetric. Their side chain could be turned around the C α -bond by 180° (Gamow et al. 1956),¹⁰ which would allow them to bind to both cavities. So,

⁹ I thank a reviewer for emphasising this contrast.

¹⁰ Gamow's explicit claim was restricted to the twelve cavities that had distinct top and bottom bases (marked * in Fig. 1), but it equally applies to the other eight cavities. He may have emphasised the 12 cavities because swapping the complementary bases in these cavities resulted in changing the combination

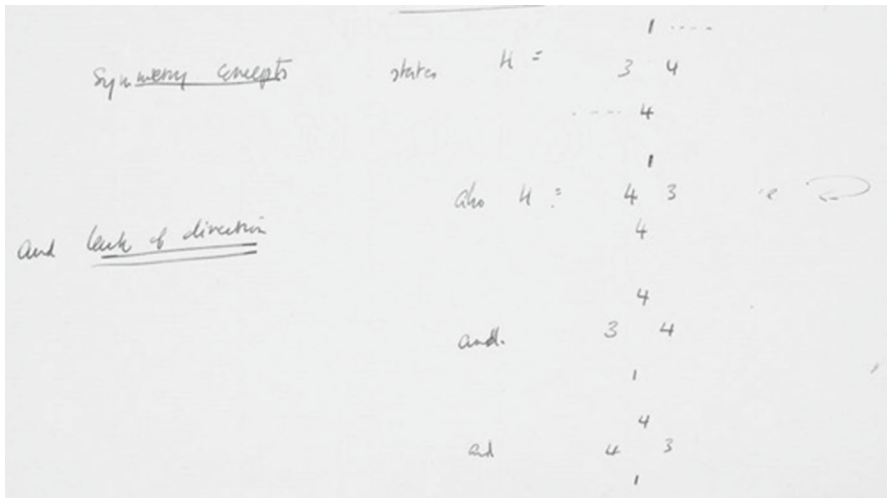


Fig. 3 Rotating the H-diamond sideways (circular arrow on the right) and bottom-up. Extract of Crick's notes on "Errors in Gamow's paper" (Crick 1950–54, p. 184 of PDF). Published with the permission of the Wellcome Collection

the first and second diamond in Fig. 3 represented cavities distinct in structure but not function (hence denoted as the same diamond H). The second difference concerned the orientation of an entire cavity within a chain of cavities. On Gamow's scheme, a cavity sitting in one direction could just as well face in the opposite direction, and without affecting its function. For example, the upper right side of the cavity depicted by the H diamond was formed by numerals 1 and 4, and its lower left side by 3 and 4 (Figs. 1 and 3). But this cavity could be positioned in the opposite direction, so that the 1–4 side came to lie on the lower left and the 3–4 side on the upper right (fourth diamond in Fig. 3). As Crick (1955) noted, Gamow simply assumed that the orientation of a cavity within a chain made no functional difference. So, diamonds one and four in Fig. 3 depicted cavities distinct in structure but not function, as did diamonds two and three. The four diamonds depicted four structurally distinct cavities that all specified the same amino acid. Indeed, H was one of 12 diamonds with four functionally equivalent variants each, marked * in Gamow's (1954a) table. The remaining 8 diamonds had two variants each.

Gamow's additional assumption (that some structurally distinct cavities were functionally equivalent) complicated the translation of known amino acid sequences into corresponding diamond chains. It raised peculiar challenges, such as deciding whether or not two distinct diamonds represented functionally equivalent cavities and demonstrating that two structurally distinct diamonds did in fact represent functionally equivalent cavities. The latter challenge arose with respect to the hypothetical leucine-specifying cavities for insulin B. Crick represented the leucine-specifying cavity that lay between the alanine- and tyrosine-specifying cavities as the T -diamond from

Footnote 10 continued

of bases lying on the same strand. Note also that the 12 cavities were initially referred to as "bilaterally symmetric" (Gamow 1954b) and later as "bilaterally antisymmetric" (Gamow et al. 1956).

Gamow's (1954b) table. That T -diamond had a 3 in its left corner (call this diamond T_3). By contrast, the diamond representing the leucine-specifying cavity above the tyrosine-specifying cavity had to have a 4 in its left corner (T_4). The difference in numerals between T_3 and T_4 implied a structural difference between the corresponding cavities, a difference that had to be functionally irrelevant since both cavities were supposed to specify leucine. So, the task was to demonstrate that, indeed, the cavities' structural difference did not affect their function.

The task could be approached by explicit reasoning in physico-chemical terms. With respect to the example above, the first step was noticing that the 3 and 4 in the corners stood for complementary bases and that therefore the two diamonds represented cavities in which the complementary bases sat on different DNA strands. The second step was to recall Gamow's assumption that this difference did not matter functionally, as long as the specified amino acid was bilaterally symmetric. Leucine happened to be bilaterally symmetric and it could therefore bind to both cavities, simply by turning its residue 180° around its $C\alpha$ -bond. Such reasoning, i.e. applying general chemical facts and theoretical assumptions to a particular case, allowed concluding that the two cavities were indeed functionally equivalent.

2.2.2 Mental rotations

Another approach, however, was much simpler: mentally rotating one diamond (T_3) and then comparing the rotated diamond with the second diamond (T_4). If they were identical, then one could infer that the two diamonds depicted the same cavity. Remarkably, a physical record of this procedure has survived in Crick's notes: the curved arrow on the right side of Fig. 3 (Crick 1950–1954, p. 184 of PDF). The arrow stood for the activity of mentally rotating the top diamond 180° sideways around its vertical axis. This manipulation moved the 3 in the left corner to the right and the 4 from right to left. Crick recorded the outcome of the rotation as the second diamond (Fig. 3). Notice that he counted both diamonds as variants of the H -diamond. Applied to the diamonds above: one rotates T_3 180° sideways and compares the outcome to T_4 . Since the rotated diamond was identical with T_4 (Fig. 3), they represented the same functional cavity, i.e. the one specifying leucine.

The diagrammatic method allowed inferring from facts about the post-rotational identities of diagrams to the functional equivalence of real-world entities. In order to see what justified these inferences, let us return to Crick's figure and consider the words "symmetry concepts" written next to the first diamond (Fig. 3, top left). The words invoked Gamow's assumption that cavities differing only in the position of their complementary base pair specify the same amino acid (due to the bilateral symmetry of the latter). Now, two cavities differing only in the position of their complementary base pair were depicted by diamonds differing only in the position of their left-right numerals (e.g. T_3 and T_4). When one such diamond was rotated 180° sideways, it became identical with the other. The post-rotational identity of diamonds therefore showed that the diamonds represented two cavities whose structural difference Gamow deemed functionally irrelevant.

Overall, we see that mentally rotating physical marks on paper generated the same conclusion as explicit reasoning with chemical facts and theoretical assumptions, albeit

with less effort. There was no need to interpret the numerals semantically, no need to remember Gamow's assumption about complementary base pairs, and no reason to verify the bilateral symmetry of particular amino acids. The diagrammatic procedure bore a striking similarity to the mental rotation tasks used in contemporary psychology in order to study visuospatial reasoning (Tversky 2005). Visuospatial reasoning is, broadly speaking, a form of reasoning about the visual and spatial properties of objects. In many rotation tasks, subjects are given two images of geometrical objects and asked to decide whether or not the objects are the same. Solving this task requires mentally rotating one of the objects until it can be compared to the other, making the comparison, and then judging whether or not the objects are the same. These are the key activities of visuospatial reasoning that we also find in the diagrammatic procedure.

So far, I focused on rotating diamonds sideways. A closely related, second type of mental transformation was rotating a diamond *upside down* around its horizontal axis by 180° . If this operation is applied to the top diamond in Fig. 3, then the 1 in its top corner moves to the bottom, and the 4 in its bottom corner ends up on top. This is the outcome Crick recorded as the third diamond (Fig. 3). Although I found no explicit record of the upside-down rotation in Crick's notes, there is no reason why Crick would have switched to explicit chemical reasoning at this point. Crick wrote "lack of direction" next to the third diamond, referring to Gamow's assumption that changing a cavity's direction within a chain was functionally irrelevant. Cavities differing only in their direction were depicted by diamonds that differed in the position of both their left-right numerals *and* their top and bottom numerals. As we saw, an example of such a pair of diamonds was the first and fourth diamond in Fig. 3, whose left-right and top-bottom numerals are reversed. Another such pair is the second and third diamond. Turning one diamond (of a pair) sideways and upside down generates an outcome that is identical to the second diamond of the pair. Again, post-rotational identity of diamonds demonstrated that distinct diamonds represented cavities with a structural difference that had no functional impact.

From these considerations it is possible to derive a simple diagrammatic procedure for determining which structurally distinct diamonds represented functionally identical cavities and which did not: all and only the diamonds that can be transformed into one another by sideways and/or upside-down rotation represented the same cavities. So, if two distinct diamonds could be transformed into one another by sideways and/or upside-down rotation, then they depicted the same cavities. On the other hand, if they could not be transformed, then they represented functionally distinct cavities. Although neither Crick nor Gamow articulated such a rule explicitly, it guided their work with diagrams. It underpinned, for instance, Crick's claim that the four diamonds in Fig. 3 are all variants of *H*. It can also be seen in the way Gamow approached a more complex task, to which I turn now.

Gamow et al. (1956) noticed that some of the known proteins had identical sequences except for one amino acid. Generating this difference with DNA cavities required changing the function of one cavity in the sequence while retaining all others. It was not obvious whether Gamow's scheme would allow for this. After all, neighbouring cavities shared bases, and so any functional change to the bases of one cavity necessitated at least structural changes in both its neighbours, which in turn had effects on their neighbours, and so on. In order to determine whether the function




Fig. 4 Rotating the *T*- and *S*-diamonds (Gamow et al. 1956, Fig. 10). Reprinted with the permission of Elsevier

of a single cavity could be altered, Gamow began by rotating the bottom diamond in a chain (Fig. 4), noting the consequences for its upper neighbour, and then repeating this process for the remaining diamonds in the chain, gradually working upwards. Rotating the *T*-diamond upside-down generated a second variant of the *T*-diamond. It also changed the left corner of its upper neighbour to a 4, and so the question arose whether that neighbour could remain an *S*-diamond. It could, by rotating the *S*-diamond sideways. This rotation changed the bottom corner of *S*'s upper neighbour to a 3. Again, the question was whether the upper neighbour could remain, in this case, a *G*-diamond. It could not, since rotating the *G*-diamond upside-down would place a 2 in its bottom corner, not a 3. So, the upper neighbour was no longer a *N* instead, and therefore represented a cavity for a different amino acid. With the *N* not rotated, the side that *N* shared with its neighbour *A* remained unaltered. Therefore, the whole *A*-diamond was not changed either. All in all, Gamow's scheme did allow changing the function of only one cavity in a chain. Or, as Gamow put it, "by rotating the diamonds, one can change one single amino acid ... without changing the rest of the sequence." Gamow's remark of "rotating the diamonds" is important here, because it shows that the authors used mental rotations to derive their conclusion.

Let us review the key findings of Sect. 2: modelling with diamond diagrams included a lot of work with physical marks on paper, and such work was clearly beneficial. Physical diagrams freed up working memory by acting as external memory stores (Sect. 2.1) and helped to replace explicit reasoning in chemical terms with visuospatial reasoning (Sect. 2.2). It is hard to see how the results, the elaboration and refutation of Gamow's hypothesis, could have been achieved without diagrams functioning as external, physical entities. These findings do not undermine the importance of imagination for scientific modelling (Toon 2016). Indeed, the mental rotations described in Sect. 2.2.2 underscore the centrality of imagination as well as the intricate to-ing and fro-ing between imagining and inscribing.

3 Is modelling an instance of distributed cognition?

The previous section described the benefits of employing physical entities (diagrams) for scientific modelling. These benefits vindicate the emphasis on external resources for our understanding of actual modelling practice (e.g. Giere 2002; Knuutila 2011; Kuorikoski and Ylikoski 2015). Some scholars go a step further and associate the importance of external resources with “distributed” or “extended cognition”. Although fleshed out in different ways, the basic idea is that scientific modelling should be construed as an instance of distributed cognition (Giere 2002; Nersessian 2006; Kuorikoski and Ylikoski 2015).¹¹ Let us take a closer look.

Giere maintains with respect to diagrams that “[w]hat is interesting about diagrammatic reasoning is the *interaction* between the diagram and a human with a fundamentally pattern-matching brain. Rather than locating all the cognition in the human brain, one locates it in the system consisting of a human together with a diagram” (Giere 2002, p. 232, italics in original). Here, the cognitive aspects of diagrammatic reasoning include not only brain states but also diagrams, i.e. things outside one’s head. Similarly, Nersessian (2006, pp. 701–702) defines ‘cognition’ in the context of model-based reasoning as comprising “a complex system, “stretched over” what have been thought of as “internal” and “external” representations and processes. [...] Cognition involves minds but extends beyond human biological capacities to encompass material artifacts and social interactions.” So, Giere (2002) and Nersessian (2006) appear to endorse the following claim¹²: a cognitive system engaged in scientific modelling comprises not only the things usually associated with cognition, i.e. brain processes, but it also includes (or can profitably be construed as including) the external, physical aspects of a scientific model. Once the reliance of scientific modelling on external resources is apparent, it is tempting to conclude that these resources are integral parts of the cognitive processes involved in modelling. We might reason as follows. The modelling activities of scientists rely not only on their internal brain processes but also, and essentially, on physical features (external models). Since the reasoning depends essentially on the physical features, these features qualify as components of cognition as much as any other essential features, such as internal brain processes.

However, it is important to realise that construing (physical) scientific models as cognitive components amounts to a demanding reading of distributed cognition and that, furthermore, a weaker reading is available. After all, philosophers of mind distinguish between two versions of the extended mind/cognition thesis (Rupert 2004;

¹¹ According to Giere (2002), we should think “of many models as being components of *distributed cognitive systems*” (p. 227, italics in original). Nersessian (2006, p. 708) holds that “model-based reasoning can be construed as occurring in distributed cognitive systems”. Similarly, Kuorikoski and Ylikoski (2015, p. 3818) “approach models as material objects that can be used to keep score of inferential moves and in this way be used to help in reasoning about the phenomena of interest. This amounts to approaching modeling as extended cognition.” Indeed, the “outsourcing of inferential work to external representational aids is a paradigmatic example of extended cognition” (p. 3823).

¹² Kuorikoski and Ylikoski (2015) are less clear on this issue, since they do not elaborate on what they call “extended cognition”. However, they explicitly associate the term with Andy Clark’s work (p. 3821), for whom cognition can extend beyond the head. This suggests that they are sympathetic to Giere’s and Nersessian’s view of distributed cognition.

Menary 2010; Wilson and Foglia 2015).¹³ According to one version, cognition depends in deep and unexpected ways on one’s environment and body, especially on external reasoning aids and one’s adeptness in handling them. The stronger version adds a further claim, which is that the external features are among the (mereological) parts of cognition. Cognition therefore literally extends from brain processes into the environment. Although both versions deviate from the traditional view of cognition, most of the controversy concerns the strong version.¹⁴ In what follows, I reserve the term “embodied cognition” for the first version and “extended cognition” for the second (following Rupert 2004), though nothing substantial hangs on this choice of words.¹⁵ What matters for present purposes is that Giere (2002) and Nersessian (2006), and perhaps also Kuorikoski and Ylikoski (2015), appear to endorse the strong, controversial variant of distributed cognition. This immediately raises the question of whether the strong version is really needed for capturing modelling with external aids or whether the weak version is in fact sufficient.¹⁶

In order to address this question, consider the theoretical roles of scientific models. For present purposes we can focus on the roles described by Kuorikoski and Ylikoski (2015). According to Kuorikoski and Ylikoski, models can enhance scientific understanding, and they can do so in at least three ways: (1) models encourage scientists to make their assumptions explicit since specifying a model’s parameters requires explicit definitions; (2) model-based inferences are more reliable than unaided inferences; (3) models expand the scope of correct what-if-inference.¹⁷ Let us grant that these are three ways in which models enhance scientific understanding and that this

¹³ Although the terms ‘mind’ and ‘cognition’ are normally used interchangeably in the extended cognition literature, this should not be understood as implying identity. Instead, cognitive processes are a subset of mental processes (e.g. Menary 2010; Wilson and Foglia 2015; but see Giere 2002).

¹⁴ It has been show, for instance, that sometimes the strong version is held for the wrong reasons: even if cognition relies essentially on external resources, this fact does not imply that they are themselves components of a cognitive process (Rupert 2004; Menary 2010; Wilson and Foglia 2015). The merit of other considerations in favour of the strong version is the subject of much debate (for reviews, see for example Menary 2010 and Wilson and Foglia 2015).

¹⁵ The terminology in this area is not standardized. For instance, Wilson and Foglia (2015) use the term “embodied cognition” for the large family of views that emphasise the role of a person’s environment and body. So, in their usage, “embodied cognition” includes, but is not limited to, what Rupert labels “embodied cognition”. I adopt Rupert’s terminology because his is one of the earliest and clearest exposition of the relation between the two varieties of extended cognition (Menary 2010; Wilson and Foglia 2015).

¹⁶ Another topic in need of clarification concerns the kinds of inferences that external aids support. Kuorikoski and Ylikoski (2015) sometimes write as if they are model-to-target inferences. For example, when considering the “outsourcing of inferential work to external representational aids”, they elaborate on “inferential work” as “the ability to use [the model] to make correct what-if inferences about the object of interest” and, in the next sentence, as drawing “inferences from the assumptions to the *explanandum*” (p. 3823, emphasis in original). Presumably, the ‘object of interest’ and the ‘*explanandum*’ refer to the model’s target, in which case the inferences at issue are model-to-target inferences. However, at least in the case of modelling with diamonds, the (primary) inferential moves aided by external diagrams were not model-to-target inferences but rather within-model inferences. That is, the externally supported inferential moves were those required for manipulating the model (diagram) itself. The moves included, for instance, the inference from a numeral in one horizontal corner to the numeral in the other, as well as inferences from mental rotations to claims about whether or not structurally distinct diamonds were variants the same type. Arguably, the external support for within-model inferences is what rendered reliable the subsequent model-to-target inferences.

¹⁷ Presumably because a model allows fine-grained tinkering with its parameters.

enhancement is a key role of models. The question then is this: in order to promote understanding in these ways, does model-based reasoning require external resources or is it also necessary that the resources are components of the cognitive process of modelling (or that they are construed as such)? If it does not matter whether or not the resources count as cognitive parts, then this suggests that the weaker form of distributed cognition is sufficient for describing modelling with external aids. Now consider the three ways of enhancing scientific understanding.

Explicitness In the case of diamond diagrams, the assumptions of Gamow's hypothesis were expressed physically by shaping external marks in certain ways and giving them a meaning. For instance, the assumption that neighbouring DNA cavities share bases was expressed by drawing adjacent diamonds so that they shared numerals and by stipulating that numerals represent bases (and diamonds depict cavities). Could this assumption have been expressed more clearly by construing the physical marks as components of a cognitive process? It is hard to see how it could have. There simply is no discernible link between the factors determining explicitness and the cognitive nature (or otherwise) of the physical marks. Consider a second example. Gamow's assumption that each cavity has one complementary base pair was expressed by limiting the numerals of left-right corners to four combinations and stipulating that the latter represent the four possible base pairs. Once these conventions were in place, Gamow's assumption was made fully explicit. There was neither a need for, nor a benefit from, construing the numerals as cognitive components.¹⁸

Reliability Identifying new diamond-amino acid assignments involved a decision about whether adding a numeral to an as-yet-unassigned diamond would render that diamond identical to a diamond that was already fully labelled. This decision required, in turn, a comparison between the diamond in question and the already labelled diamonds in the chain. Recording the labelled diamonds on paper generated a lasting, physical resource for such comparisons. By contrast, there is no indication that the comparisons could have been made more reliable by counting the physical records as parts of cognition. Construing the records as parts of cognition, or failing to do so, has no detectable effect on the reliability of modelling with diamond diagrams.

Scope The diamond diagrams did not increase the scope of *correct* what-if inferences because Gamow's hypothesis was false. Nevertheless, we can ask: (1) Would the diagrams have increased the scope if the hypothesis had been correct? (2) Would construing the diagrams as cognitive components have affected the scope? In order to answer these questions, recall Gamow's attempt to determine whether his proposed mechanism allowed exchanging exactly one amino acid in a protein (Sect. 2.2.2, Fig. 4). It is easily seen that Gamow's attempt generated a series of what-if inferences, from inferences about diagrams and DNA cavities to inferences about amino acids. For example, if the complementary bases of the cavity represented by the *S*-diamond were swapped, then a different amino acid would be added. And if Gamow's hypothesis had been true, this would have been a correct what-if inference.¹⁹ Since such examples can

¹⁸ Making assumptions diagrammatically explicit was crucial for modelling purposes. For instance, if adjacent diamonds did not share numerals, then it would be impossible to derive the numerals of one diamond from those of the other.

¹⁹ The same goes for the following inference: If the complementary bases of the cavity represented by *S* were swapped, then the functional identity of its upper neighbour would change.

be multiplied, it is safe to conclude that the diagrams would have allowed expanding the scope of correct what-if inferences. Would the scope have been affected by construing the diagrams as cognitive components? Again, it is hard to see how it could have. Deriving what-if inferences was the result of modifying the external features of a symbol system and tracing the consequences. This capacity was not affected by regarding the external features as cognitive parts.

Overall, it can be seen that the diagrammatic method improved scientific understanding by increasing the reliability of reasoning and by making theoretical assumptions explicit. Furthermore, these achievements can be explained by the fact that the method relied on external, physical marks. There is no evidence that construing the marks as parts of cognition makes a difference. These findings strongly suggest that embodied cognition is sufficient for explaining how the diagrammatic method improved scientific understanding. That said, the findings do not exclude the possibility that construing the marks as cognitive might affect other kinds of theoretical roles or that it might deliver benefits of an altogether different sort. However, at present these are theoretical possibilities without positive support. So, on balance, the current evidence favours the conclusion that the diagrammatic method is best understood as an instance of embodied, rather than extended, cognition (in Rupert's 2004 sense).

The fact that modelling with diamond diagrams was very likely an instance of embodied cognition raises the prospect that other kinds of physical models also rely on embodied cognition (instead of extended cognition). Of course, diamond diagrams may be the exception among physical models; here we reach the limitations of a case study. Nevertheless, the present findings provide good reasons for focusing attention on embodied cognition and investigating the scope of this less controversial form of distributed cognition.²⁰

4 Conclusion

This paper investigated how the physical features of scientific models affect their use and whether modelling is extended cognition. These topics were approached by means of a historical case study in which diagrams served as models for elaborating a mechanistic hypothesis.

I argued that scientists employed the diagrams (physical models) as external information stores to which they returned repeatedly in order to perform subsequent modelling steps. More specifically, external diagrams were instrumental for decomposing complex tasks into many smaller steps so that the work with diagrams alternated between mental manipulations, recording their results, and then feeding the results back into subsequent mental manipulations. In addition, physical models allowed replacing explicit reasoning about chemical facts and theoretical assumptions with the visuospatial manipulations of diagrams, which are cognitively less demanding. Both Crick's arrow (Fig. 3) and Gamow's remark about "rotating the diamonds" provide historical evidence that visuospatial reasoning in the form of mentally rotating

²⁰ I thank a reviewer for prompting me to clarify the implications of this case study for distributed cognition.

diamonds was part of modelling with diagrams. Therefore, the models' physicality affected the very nature of scientific reasoning.

This study also explored the link between distributed cognition and modelling with physical aids. I adopted the distinction between embodied and extended cognition (both forms of cognition rely on physical marks, but extended cognition demands in addition that the marks are cognitive components) and explored how the diagrammatic method advanced scientific understanding in the ways described by Kuorikoski and Ylikoski (2015). It can be shown that the method's reliance on external marks was critical, whereas construing the marks as cognitive was not. Consequently, the weaker and less controversial form of distributed cognition is sufficient for understanding modelling with diamond diagrams.

Acknowledgements I thank two anonymous reviewers for helpful objections and suggestions. I gratefully acknowledge the permissions granted by Elsevier, The Royal Danish Academy of Sciences and Letters, and the Wellcome Collection to reproduce copyrighted material.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Craver, C. (2008). Physical law and mechanistic explanation in the Hodgkin and Huxley model of the action potential. *Philosophy of Science*, 75(5), 1022–1033.
- Craver, C., & Darden, L. (2013). *In search of mechanisms: Discoveries across the life sciences*. Chicago: University of Chicago Press.
- Crick, F. H. (1950–54). *Errors in Gamow's code*. PP/CRI/H/1/20/1:Box 70. <http://wellcomelibrary.org/player/b18181600> (Unpublished notes). Accessed 15 Aug 2018.
- Crick, F. H. (1953–1954). *Gamow's code*. PP/CRI/H/1/20/2:Box 70. <http://library.wellcome.ac.uk/player/b18187328> (Unpublished notes). Accessed 15 Aug 2018.
- Crick, F. H. (1955). *On degenerate templates and the adaptor hypothesis: A note for the RNA tie club*. January 1955. PPCRI/H/1/38. <http://library.wellcome.ac.uk/player/b18186300> (Unpublished note). Accessed 15 Aug 2018.
- Frigg, R., & Hartmann, S. (2012). Models in science. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/spr2006/entries/models-science/>. Accessed 1 May 2018.
- Gamow, G. (1954a). Possible relation between deoxyribonucleic acid and protein structures. *Nature*, 173, 318.
- Gamow, G. (1954b). Possible mathematical relation between deoxyribonucleic acid and proteins. *Det Kongelige Danske Videnskabernes Selskab, Biologiske Meddelelser*, 22(3), 3–13.
- Gamow, G., Rich, A., & Yčas, M. (1956). The problem of information transfer from the nucleic acids to proteins. *Advances in Biological and Medical Physics*, 4, 23–68.
- Giere, R. N. (2002). Models as parts of distributed cognitive systems. In L. Magnani & N. J. Nersessian (Eds.), *Model-based reasoning: Science, technology, values* (pp. 227–241). New York: Kluwer Academic/Plenum Publishers.

- Goodwin, W. (2008). Structural formulas and explanation in organic chemistry. *Foundations of Chemistry*, 10, 117–127.
- Hegarty, M., Stieff, M., & Dixon, B. (2015). Reasoning with diagrams: Towards a broad ontology of spatial thinking strategies. In D. R. Montello, K. E. Grossner, & D. G. Janelle (Eds.), *Space in mind: Concepts for spatial learning and education* (pp. 75–98). Cambridge, MA: MIT Press.
- Hegarty, M., & Stull, A. T. (2012). Visuospatial thinking. In K. J. Holyoak & R. G. Morrison (Eds.), *The oxford handbook of thinking and reasoning* (pp. 606–630). Oxford: Oxford University Press.
- Knuuttila, T. (2011). Modelling and representing: An artefactual approach to model-based representation. *Studies in History and Philosophy of Science*, 42, 262–271.
- Kuorikoski, J., & Ylikoski, P. (2015). External representations and scientific understanding. *Synthese*, 192(12), 3817–3837.
- Menary, R. (2010). Introduction: The extended mind in focus. In R. Menary (Ed.), *The extended mind* (pp. 1–26). Cambridge, MA: MIT Press.
- Nersessian, N. J. (2006). Model-based reasoning in distributed cognitive systems. *Philosophy of Science*, 73(5), 699–709.
- Perini, L. (2012). Truth-bearers or truth-makers? *Spontaneous Generations: A Journal for the History and Philosophy of Science*, 6(1), 142–147.
- Rupert, R. D. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, 101(8), 389–428.
- Sanger, F., & Tuppy, H. (1951). The amino-acid sequence in the phenylalanyl chain of insulin. 2. The investigation of peptides from enzymic hydrolysates. *Biochemical Journal*, 49, 481–490.
- Toon, A. (2016). Imagination in scientific modeling. In A. Kind (Ed.), *The Routledge handbook of philosophy of imagination* (pp. 451–462). Oxford: Routledge.
- Tversky, B. (2005). Visuospatial reasoning. In K. J. Holyoak & R. G. Morrison (Eds.), *The Cambridge handbook of thinking and reasoning* (pp. 209–240). Cambridge: Cambridge University Press.
- Wilson, R. A., & Foglia, L. (2015). Embodied cognition. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.Stanford.edu/archives/spr2017/entries/embodied-Cognition/>. Accessed 1 May 2015.