

Constructing Diagrams to Understand Phenomena and Mechanisms

Benjamin Sheredos
William Bechtel

Department of Philosophy
& Center for Circadian Biology
UC San Diego

Biologists often hypothesize mechanisms to explain phenomena. Our interest is how their understanding of the phenomena and mechanisms develops as they construct diagrams to communicate their claims. We present two case studies in which scientists integrate various data to create a single diagram to communicate their major conclusions in a research publication. In both cases, the history of revisions suggests that scientists' initial drafts encode biases and oversights that are only gradually overcome through prolonged, reflective re-design. To account for this, we suggest that scientists only develop a unitary understanding of their results *through* their attempts to communicate them.

1. Introduction

In biology, explanation often involves characterizing a phenomenon and generating an account of the *mechanism* thought to be responsible for it. The notion of mechanism has played this role in the life sciences since at least the 18th century, when it was adopted to characterize explanations that result from analyzing or decomposing biological systems into component parts, detailing their operations, and determining how these parts are organized and the operations orchestrated to produce the phenomena of interest (Bechtel & Richardson, 1993/2010). Scientists frequently find it productive to represent both phenomena and mechanisms in diagrams, in which different glyphs (Tversky, 2011) are laid out spatially. Shapes represent entities (the mechanism or its parts) and arrows represent operations. Space-on-the-page sometimes represents physical space (e.g., the nucleus versus the cytoplasm of the cell) but often is used simply to separate glyphs, distinguishing the represented parts and operations (Sheredos, Burnston, Abrahamsen, & Bechtel, 2013). Often, viewers can mentally animate a diagram to get an intuitive understanding of the mechanism's operations (Hegarty, 1992). Diagrams also aid in producing abstract mathematical cognition in the construction of computational models (Jones & Wolkenhauer, 2012).

Our focus here is on how scientists *generate* such diagrams. Generally these figures do not arise in a final format all at once, but result from a history of producing and revising interim drafts. Hand-drawn sketches might be preserved in laboratory notebooks (Nersessian, 2008). In the electronic era, a lineage of drafts is often preserved digitally in the files researchers save on the way to a final diagram. We take advantage of these to study the development of diagrams that appeared in two published papers. The first authors of each paper have provided their drafts of both text and the figures, allowing us to analyze their development. The last figure in one paper depicts a mechanism proposed to explain a phenomenon, whereas in the other the last figure presents a new phenomenon to be explained. We examine the drafts leading to these.

The researchers made major revisions to these diagrams as work on the manuscripts proceeded, reflecting the cognitive labor required in their development. We advance a hypothesis regarding why such labor is required, and why it takes the form it does: prior to the attempt to develop a coherent diagram, scientists themselves lack full understanding of the domain of inquiry. They have a “cognitive collage” (Tversky, 1993)

consisting of somewhat isolated and only partly-integrated understandings, often gleaned from diverse sources of data. It is through the attempt to develop a communicative diagram that these disparate, partial understandings are integrated into a detailed, cohesive whole.

Our case studies both concern research on circadian rhythms: endogenously-generated oscillations of approximately 24 hours that regulate the timing of other physiological and behavioral activities. The laboratory from which these publications arose studies circadian rhythms in cyanobacteria (specifically *Synechococcus elongatus*), the only bacterial lineage in which circadian rhythms have been demonstrated.

The basic mechanism responsible for circadian timekeeping in cyanobacteria is represented in Figure 1. (This figure comes from our first case study, examined further in the next section. For further details on the core mechanism see Kim, Dong, Carruthers, Golden, & LiWang, 2008.) The mechanism involves three proteins (KaiA, KaiB, and KaiC) plus their states and interactions at four major time-points (organized here in a circle, with different time-points at top, right, bottom, and left). KaiC is the large macromolecule shown at each time-point. It undergoes phosphorylation and dephosphorylation at two locations; the added phosphate groups are symbolized by the letter P in a black circle. KaiC itself initiates both phosphorylation and dephosphorylation, but the other two Kai proteins determine which dominates. When KaiA, represented using a purple, “bunny-eared” icon (top, right, and bottom), binds to KaiC (see top) phosphorylation is sped up and KaiC quickly becomes phosphorylated at both locations (see the two “P”s at right). When KaiB, represented using four stacked red ovals (at right and bottom) binds, it sequesters KaiA (see bottom), allowing dephosphorylation to proceed until neither location is phosphorylated (see left).

Since there is a specific and regular order of phosphorylation, and one cycle takes about 24 hours, KaiC’s phosphorylation state predicts the current time of day, and serves as the cyanobacterium’s “clock.” Although open questions remain, this basic mechanism is well-established (see Mackey, Golden, & Ditty, 2011, for review) and provides the backdrop for the research pursued in our two case studies.




Figure 1. The first figure in Paddock et al. (2013) showing KaiC’s phosphorylation cycle, as regulated by KaiA&B. Arrows leading to the word *Output?* encode uncertainty regarding which phosphorylation state communicates predicted time-of-day to the rest of the cell.

2. Advancing a New Hypothesis About the Output Mechanism

In our first case, Paddock, Boyd, Adin, and Golden (2013) advanced an important revision of what had become the standard account of the output mechanism through which the cyanobacterial circadian clock regulates the expression of virtually all genes. Two relatively well-defined classes of genes exhibit peak expression around (predicted) dawn and (predicted) dusk. If expression peaks near dawn, the gene is said to be regulated by a Class 1 promoter; if near dusk, by a Class 2 promoter. The proximal cause at work in each case is a transcription factor, which activates the promoter and initiates gene expression. Somehow, the clock must regulate promoter activation. Yet the Kai proteins are not transcription factors: none can directly influence any gene's expression. So additional components, forming an "output pathway," must mediate clock control of gene expression.

Two proteins, SasA and RpaA, had long been implicated since knocking out these proteins severely reduces rhythmic gene expression, even though the clock (KaiC's phosphorylation rhythm) is left intact. Since RpaA, but not SasA, is a transcription factor, the output pathway was hypothesized to run from KaiC to SasA to RpaA (Takai, Nakajima, Oyama, Kito, Sugita, Sugita, Kondo, & Iwasaki, 2006). We return to discuss this SasA-RpaA pathway below. What this research had not been able to determine, however, was which phosphorylation state(s) of KaiC triggers output.

This question is posed in Figure 1 (which is Paddock et al.'s published Figure 1). In addition to the glyphs we discussed above, the graphic includes four dotted arrows, originating at each phosphorylation state of KaiC and terminating in "Output?" Use of question marks to indicate uncertainties is common in mechanism diagrams. These arrows were added in a late draft (March 13, 2013), but the uncertainties they represent were formulated well in advance, as the specific target of research: Paddock et al. tested which phosphorylation state(s) drives output from the clock, and controls gene expression. The decisive experiments involved two steps.

First, Paddock et al. took cells and knocked out KaiC, destroying the clock. In this condition, there is no circadian *regulation* of gene expression: transcription factors activate promoters at their leisure. It was observed that with circadian regulation eradicated, Class 1 promoters "default" to a constantly high level of activation compared to wild-type (rather than selectively increasing activation at dawn) and Class 2 promoters "default" to constantly low activation (rather than selectively increasing activation at dusk).

Next, Paddock et al. reasoned that any phosphorylation state of KaiC that induced a deflection *away* from these "default" values could play some role in controlling output. To examine this, they created four molecules, each of which mimicked one phosphorylation state. (These phosphomimetics are named in italics in Figure 1). They then replaced KaiC with one of the phosphomimetics. Each modified cell essentially has a clock that is artificially "stopped" at one time-of-day. Paddock et al. then measured the effect on gene expression (using a luciferase reporter to detect promoter activation).

Only one phosphomimetic (*KaiC-ET*) induced activation different from the KaiC knockout's "default" activation. It both repressed the default-high activation of Class 1 promoters and enhanced the default-low activation of Class 2 promoters. This established that a single phosphorylation state of KaiC serves as the clock's output signal. This phosphorylation state (labeled "KaiC-pST" on the bottom of Fig. 1) corresponds, roughly, to predicted middle-of-the-night.

Between October 2012, when the authors began writing the manuscript, and June 2013, when they submitted it for publication, Paddock et al. drafted a series of diagrams to resolve the uncertainties presented in Figure 1. Two early versions are shown in Figure 2. Figure 2A continues to show all four phosphorylation states of KaiC, and adds an inhibitory arrow showing repression of P_{kaiBC} (a Class 1 promoter which serves to represent all Class 1 promoters) and an excitatory arrow showing activation of P_{purf} (representative of Class 2 promoters). Figure 2B partially simplifies the diagram by leaving out phosphorylation states that were ineffective in regulating gene expression, retaining a circle to indicate the phosphorylation cycle of KaiC. It also adds some linguistic labels and an indication that the clock is affected by inputs.

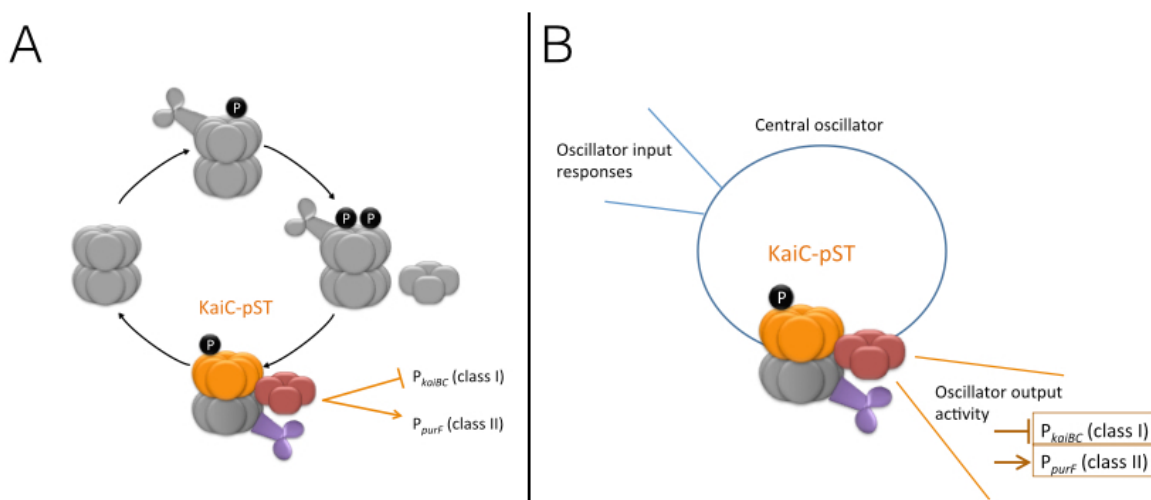


Figure 2. Panel A: an early version in the lineage of sketches that culminated in Figure 7, dated December 4, 2012. Panel B: a pared-down version, dated January 11, 2013.

These early drafts foregrounded the importance of KaiC-pST in regulating clock output, but did not include roles for SasA and RpaA, which were known to influence output. Paddock et al.'s data showed that when their phosphomimetic induced output from the clock, it did not affect the SasA-RpaA pathway. This left a puzzle: RpaA had been supposed to *mediate* the output, and yet it was not affected by the newly-identified output signal. To resolve this puzzle, the researchers drew upon additional data involving RpaA knockouts. As noted above, with KaiC knocked out (but with RpaA present), Class 1 promoters “default” to high activation, and Class 2 promoters “default” to low activation. In a RpaA knockout with KaiC still present these values are reversed: Class 1 promoters show constant low activation, and Class 2 promoters show constant high activation. When both RpaA and KaiC were knocked out, the results match those observed in the KaiC knockout alone. Taken together, the data indicate (a) that KaiC-pST affects output independently of RpaA, and (b) that the influence of RpaA is antagonistic to, or inhibitory of, the effects of KaiC-pST. Paddock et al. concluded that there were *two* output pathways: the previously-known SasA-RpaA pathway, and the one demonstrated to originate from KaiC-pST.

They drafted new diagrams (starting March 2013) to try to show how the two pathways interact in regulating gene expression. Presumably because they wanted to show the origin of the SasA-RpaA pathway in a different phosphorylation state, they restored the

other phosphorylation states that had been dropped from Figure 2B (see Figure 3A below). Because the origin of the SasA-RpaA pathway was unknown, the sketch shown in Figure 3A does not link it to any one phosphorylation state. It is shown as inhibiting output from the KaiC-ST phosphoform. Here effects on gene expression are shown all at once, in terms of a general measure, Oscillator Output Activity, which we do not discuss further (an analysis has been provided by Burnston, Sheredos, Abrahamsen, and Bechtel, 2014).

Altogether Paddock et al. generated seven variants of Figure 3A. These became quite complex as they tried to illustrate the interactions between pathways. Then in the draft of April 11 they abruptly changed to the simpler format shown in Figure 3B. Multiple representations of KaiC's phosphorylation states are removed, and a single circle represents KaiC's phosphorylation rhythm. Another circle represents a cycle of RpaA phosphorylation. Instead of trying to show the effects on Class 1 and Class 2 promoters at once, they show each separately, duplicating the whole arrangement. The schematic graphs on the right indicate effects on expression (using recorded bioluminescence as well as the measure of Oscillator Output Activity, now renamed Kai Oscillator Activity (KOA)).

Figure 3B shows what appeared, with minor changes, as part of Paddock et al.'s Figure 7. The history of drafting described here reflects a variety of attempts to portray the mechanisms of circadian output. What we highlight is the progression through several repetitive phases of abstraction, or the elimination of detail. In moving from Figure 2A to 2B, a number of details regarding phosphorylation states of KaiC are deemed irrelevant and dropped out. Yet when it comes time to add a depiction of the RpaA pathway, these same details re-appear in Figure 3A. Along with a number of other changes, there is a repetition of the same abstraction to obtain 3B, and the same details are again dropped out. This months-long drafting process only gradually produced the published figure, and one sees the researchers struggling repeatedly to move away from their initial, detail-rich sketch.

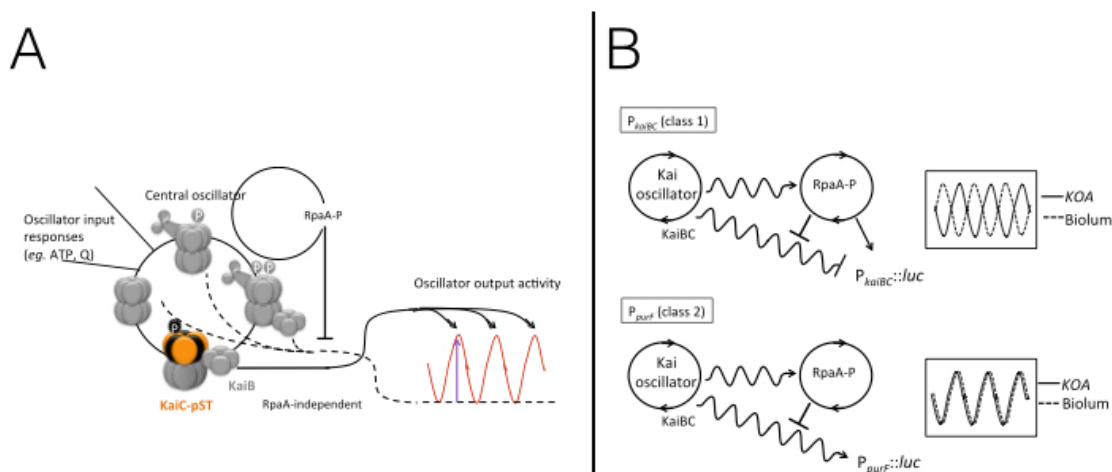


Figure 3. Panel A: one of several intermediate versions that appeared in versions of the manuscript from early March until mid-April 2013. Panel B: a pared down version that appeared on April 11, 2013.

3. Characterizing the Changing Location of the Clock within the Cell

Our second case comes from the same laboratory. Instead of advancing a new set of operations in a mechanism, Cohen, Erb, Selimkhanov, Dong, Hasty, Pogliano, and Golden (2014) reveal a new phenomenal aspect of the circadian clock, its changing location within the cell over the course of a day. Although the clock's migration is potentially important in explaining the operation of the clock, the goal of the paper is simply to demonstrate this movement.

Since they lack internal membranes, bacteria were long regarded as internally disorganized bags of genes, enzymes, and other molecules. Recent research has identified extensive internal organization and determined its importance for various physiological activities of bacteria (Rudner & Losick, 2010). Cohen et al. set out to investigate where the Kai proteins are located in the cell. Using luciferase and fluorescence reporters, they determined that although KaiA and KaiC are distributed throughout the cell during the day, at night they localize to one pole. Notably, when KaiA and KaiC are localized at the cell pole, they are co-localized with KaiB and with CikA (a part of the input pathway to the clock, affecting its “entrainment” or synchronization to local day/night cycles). Cohen et al. suggest that this localization may be functionally significant for timekeeping, and may “facilitate interactions among the clock components” (p. 1840).

The data graphics for the paper, presenting evidence for the changing localization of KaiA and KaiC, were largely settled by the time drafting of the manuscript began in March 2014. Over the following four months of drafting, much effort was spent developing a diagram linking the localization of the proteins to previously-known operations involved in the clock. All versions were prepared by the first author, with others offering advice and aiding decisions between alternatives. We examine a few steps between the initial draft and the final figure, which eventually appeared as the final figure in the published article.

The initial draft, dated March 7, 2014 comprised five panels. Three panels reproduced extant images (from the web or from another publication) to present some examples of how other diagrams had shown relevant information. One panel consisted of questions and design considerations for the figure, and read:

“Model: entrainment/proteolysis.

1. SDH/respiration goes to poles in low light.
2. ATPase interactions at night? ATPases are in the curved part of the chloroplasts
3. Curvature”

Although these phrases are cryptic, they reference specific information that the researchers considered including in their diagram. The word *entrainment* refers to CikA's role in the input pathway. The co-localization of CikA and KaiC at the cell pole may be related to entrainment, and it was considered whether to emphasize this in a future draft. The word *proteolysis* invokes a well-documented migration of proteins to the membrane when targeted for destruction. The newly-documented movement of the Kai proteins was found not to be related to proteolysis, and it was considered that future sketches might underscore this. The first bullet point points out that succinate dehydrogenase (SDH) and other enzymes involved in respiration also migrate to the poles. Implicitly, the question of the relation of the clock's migration to basic cell metabolism is being raised. The next bullet point raises it more explicitly, asking whether the clock's migration is related to energetic

processes at the pole. The last bullet point raises the question of whether the curvature of the membrane at the pole figures in directing the migration.

Although these were raised as design considerations, none were addressed in the initial draft that appeared in the remaining panel (Figure 4 below). A single cyanobacterium is shown with a green line representing its membrane. To illustrate the different state of the clock over time, the figure is divided diagonally into two segments (day phase is yellow, including an icon of the sun, and night is grey, including a moon). In each half of the figure, the phosphorylation cycle is shown, using glyphs similar to Paddock et al.'s for the Kai proteins, but showing only two of the four phosphoforms (white circles indicate phosphates). During the day phase KaiA is shown bound to phosphorylated KaiC, and detached from unphosphorylated KaiC; all these glyphs are situated towards the center and away from the pole. In the night phase CikA and KaiB are included, and the glyphs are placed near the pole. Two bullet points reference other studies documenting events occurring during the night phase.

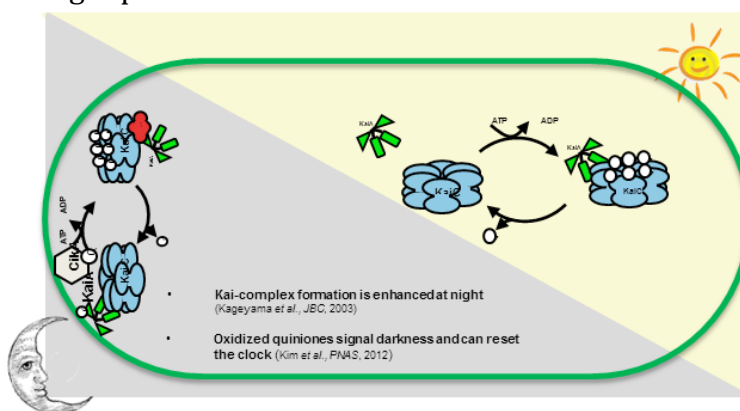


Figure 4. First draft of the mechanism diagram in Cohen et al. Dated March 7, 2014.

This first draft did not address any of the additional design issues discussed above (entrainment, proteolysis, metabolism, etc.). Based on feedback the first author received, she prepared two revised versions (Figure 5 below). Neither addressed those additional design issues but instead focused on the phenomenon of localization alone. Two features shared by these new drafts are the use of a vertical rather than diagonal division of the figure into sections for day and night, and the incorporation of a diamond representation in the center for the full, four-stage cycle of KaiC phosphorylation.




Figure 5. Two drafts of Cohen et al. Figure 6 from April 6, 2014.

The version in Figure 5's left panel retains the portrayal of one bacterial cell, with the Kai complex localized to the pole at night (now on the right) and free in the cytoplasm during the day. Overlaid is a bell-curve representation of the abundance of KaiC, which increases in concentration during the day and declines through the night. The implied x-axis for this graph imposes a linear representation of time, from dawn on the left to the next day's dawn on the right. This is in tension with the cyclical representation of time in the center of the figure. This kind of infelicity is not surprising in a draft diagram, as the author is actively trying out ideas in the attempt to construct a coherent representation. Despite this infelicity, the first author prefers this version, and continues to use it in her talks.

Other members of the research team, however, preferred the version in Figure 5's right panel, which introduces a fundamental change: multiple representations of the cell, aligned with the four phosphorylation states of KaiC, shown in the cyclical representations in the center. It is interesting that the whole figure takes the form of an oval although there is no longer any attempt to show all processes within a single bacterial cell: this is a "remnant" from the first sketch. Finally, this figure introduces a legend to link different glyphs to the molecules they represent.

After feedback from the other authors, the first author created two more versions by April 25, 2014. One of these (Figure 6 below) was eventually published without any further alteration. The sun and moon glyphs are re-introduced, and the spacing of some protein glyphs is slightly altered. Perhaps the most significant innovation is that additional KaiC icons in light blue are added in all stages of the cycle. This is intended to indicate that there are many copies of KaiC in the cell, and they are often in different states of phosphorylation.




Figure 6. Final version of Figure 6 in Cohen et al.

From the initial sketch, the diagram underwent substantial modification until the authors settled on the published version. We highlight that the final version includes much *more* detail than was present in the initial sketch, and that the history of revision is one of gradually adding details. The first step was to consider a variety of details that had been omitted (e.g., regarding entrainment, proteolysis, metabolism). These were not addressed; instead the authors added a clockwise portrayal of the progression of phosphorylation

states, and exploited this to organize the newly-discovered information regarding Kai localization at different times of day. The fundamental constraint was to deploy limited space-on-the-page to simultaneously represent intracellular space, functional states of the components, and time-of-day. This proved to be difficult: several rounds of revision were required before a format was attained which overcame the limitations of the initial sketch. Even still, many graphical elements changed little through the revisions.

4. An Hypothesis: Scientists Develop Understanding by Drafting Diagrams

We examined the evolution of two diagrams developed to communicate hypotheses about phenomena and mechanisms. Through analysis of the lineage of drafts the authors made, we identified an iterative process in which different representational strategies were gradually developed and enacted. In one case, the initial sketch was much more detailed than the final graphic; the excess details stubbornly re-appeared mid-way through revisions, only to be dropped again, revealing an iterative attempt to pare down irrelevant detail. In the other case, the initial sketch required supplementation to reach the desired degree of detail, and basic limitations of the initial sketch had to be gradually overcome in order for details to be coherently included. We described an iterative attempt to add in relevant detail.

One might regard the development of these diagrams as essentially epiphenomenal to scientific cognition: at the outset, the researchers possessed a cohesive understanding of the domain, and diagram construction was an additional layer of practice, aimed at developing a representation to aid in *communicating* that pre-established understanding. We question this “epiphenomenalist” view. The histories of revision suggest that at the outset it was neither obvious to scientists what details should be included in an adequate diagram, nor obvious how relevant details might be adequately represented. Rather, an initial attempt was made, and its excesses and omissions were *then* identified and corrected. The epiphenomenalist might account for this by proposing some general cognitive inability to communicate the cohesive understanding of the domain which researchers allegedly had in advance. But construction of these diagrams was preceded by months of reflective and careful experimental work, resulting in a hard-won understanding of the domain that motivated the researchers to write a manuscript in the first place. We can agree with the epiphenomenalist that translating the pre-established understanding into a specifically *graphical* format is an important challenge. But this is not the whole story. First, the data, which support the understanding of the domain, are typically *already* encoded in a graphical format – in the data graphics which, in our cases, were essentially finalized before authors begin the months-long process of developing their diagrams. (Moreover, both our cases show researchers integrating data-graphics directly into mechanism diagrams, suggesting there is little if any cognitive “gap” between them.) Second, while it takes multiple revisions to generate a diagram which is deemed “just right,” there is little reason to posit any inability to develop graphical representations as such: witness the number and variety of graphics the researchers developed, of which we have given only a small sample.

An adequate account must grant that researchers began with *some* understanding of the hypotheses they aim to communicate, but must account for the great expenditure of cognitive labor documented in the history of revisions. We propose that as a result of experimental work, scientists understand the domain through a variegated “cognitive

collage” (Tversky 1993) involving a diversity of representational formats. Some might be abstract (e.g., mathematical), others more clearly materially grounded (e.g., embodied familiarity with experimental protocols), and most are probably a mix. These representations are sufficiently integrated to enable the researcher to articulate their major hypothesis, but they are not yet integrated in a single representation that simultaneously provides an adequate understanding of the evidence for, and relations between, various elements of the hypothesis. There is at this point no “map-like” representation that integrates all this information. The initial sketches, we propose, are the first attempt to integrate this information into a cohesive representation. This integrative process is prone to what scientists regard as errors, of which we have identified two kinds: the inclusion of irrelevant and the omission of relevant detail. Moreover, the process is prone to a kind of anchoring effect—the initial sketch may include infelicitous elements that persist as an obstacle for later revisions (e.g., the re-appearance of irrelevant detail in Paddock et al.'s revisions; the difficulty of representing time in Cohen et al.'s graphics).

A variety of authors have argued for the practical necessity and epistemic merits of “multiple models idealization” (see Weisberg 2007 for references to the many proponents of this view). Viewed in the context of that work, our hypothesis provides an account of how and when such integration can be achieved: an iterative process of diagram redesign can generate the final, cohesive understanding of the phenomena or mechanism.

References

- Bechtel, W., & Richardson, R. C. (1993/2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Cambridge, MA: MIT Press. 1993 edition published by Princeton University Press.
- Burnston, D. C., Sheredos, B., Abrahamsen, A., & Bechtel, W. (2014). Scientists' use of diagrams in developing mechanistic explanations: A case study from chronobiology. *Pragmatics and Cognition*.
- Cohen, S. E., Erb, M. L., Selimkhanov, J., Dong, G., Hasty, J., Pogliano, J., & Golden, S. S. (2014). Dynamic localization of the cyanobacterial circadian clock proteins. *Current Biology*, *24*, 1836-1844.
- Hegarty, M. (1992). Mental animation: Inferring motion from static displays of mechanical systems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 1084-1102.
- Jones, N., & Wolkenhauer, O. (2012). Diagrams as locality aids for explanation and model construction in cell biology. *Biology and Philosophy*, *27*, 705-721.
- Kim, Y.-I., Dong, G., Carruthers, C. W., Golden, S. S., & LiWang, A. (2008). The day/night switch in KaiC, a central oscillator component of the circadian clock of cyanobacteria. *Proceedings of the National Academy of Sciences*, *105*, 12825-12830.
- Mackey, S. R., Golden, S. S., & Ditty, J. L. (2011). The itty-bitty time machine: Genetics of the cyanobacterial circadian clock. *Advances in genetics*, *74*, 13-53.
- Nersessian, N. (2008). *Creating scientific concepts*. Cambridge, MA: MIT Press.
- Paddock, M. L., Boyd, J. S., Adin, D. M., & Golden, S. S. (2013). Active output state of the *Synechococcus Kai* circadian oscillator. *Proceedings of the National Academy of Sciences*, *110*, E3849-E3857.

- Rudner, D. Z., & Losick, R. (2010). Protein subcellular localization in bacteria. *Cold Spring Harbor Perspectives on Biology*, 2, a000307.
- Sheredos, B., Burnston, D., Abrahamsen, A., & Bechtel, W. (2013). Why do biologists use so many diagrams? *Philosophy of Science*, 80, 931-944.
- Takai, N., Nakajima, M., Oyama, T., Kito, R., Sugita, C., Sugita, M., Kondo, T., & Iwasaki, H. (2006). A KaiC-associating SasA-RpaA two-component regulatory system as a major circadian timing mediator in cyanobacteria. *Proceedings of the National Academy of Sciences*, 103, 12109-12114.
- Tversky, B. (1993) "Cognitive maps, cognitive collages, and mental spatial models" *Spatial Information Theory a Theoretical Basis for GIS*, volume 716 of the series *Lecture notes in computer science*, 14-14.
- Tversky, B. (2011). Visualizing thought. *Topics in Cognitive Science*, 3, 499-535.
- Weisberg, M. (2007) Three kinds of Idealization. *J.Phil* 104(12), 639-659.