Introspection in Problem Solving

Frank Jäkel¹ and Cornell Schreiber²

¹*Institute of Cognitive Science, Universität Osnabrück*

²*Department of Philosophy, Universität Wien*

Correspondence:

Frank Jäkel, Institute of Cognitive Science, Universität Osnabrück.
Email: fjaekel@uos.de

Keywords:

introspection, problem solving, metacognition, verbal reports, think aloud

Problem solving research has encountered an impasse. Since the seminal work of Newell and Simon (1972) researchers do not seem to have made much theoretical progress (Batchelder and Alexander, 2012; Ohlsson, 2012). In this paper we argue that one factor that is holding back the field is the widespread rejection of introspection among cognitive scientists. We review evidence that introspection improves problem solving performance, sometimes dramatically. Several studies suggest that self-observation, self-monitoring, and self-reflection play a key role in developing problem solving strategies. We argue that studying these introspective processes will require researchers to systematically ask subjects to introspect. However, we document that cognitive science textbooks dismiss introspection and as a consequence introspective methods are not used in problem solving research, even when it would be appropriate. We conclude that research on problem solving would benefit from embracing introspection rather than dismissing it.

In contemporary cognitive science introspection is widely regarded as a mysterious and problematic method for uncovering the operations of the mind. Most researchers will typically avoid the conceptual and methodological minefield of introspection. Johnson-Laird (2008, p. 17), in the context of reasoning, expresses an attitude towards introspection that many cognitive scientists share: “To reason is to carry out a mental process. Introspection doesn’t reveal to us how that process works. (If it did, then psychologists would have understood how it worked long ago.) Hence the process is unconscious.” There are good reasons to be skeptical about the usefulness of introspection as a method in many areas of perception and cognition. Many processes are unconscious and hence, by definition, inaccessible for introspection. From this, however, it does not follow that introspection has no role to play in psychological research or that introspection is not an interesting process worthy of study in itself.

In the context of problem solving it is important to stress that introspection is, at least sometimes, an essential part of problem solving processes. When dealing with new or hard problems that cannot be solved by standard methods, good problem solvers introspect while they engage with their task: they examine their strategies and representations, they evaluate their progress, and they realize that they are stuck or that they have had an insight. Understanding often seems to depend on some form of conscious and deliberate reflection on one’s own thoughts. These are salient phenomena that we should not ignore. In the literature, the dubious term introspection is often avoided and is replaced with notions such as self-observation, self-monitoring, self-reflection, or metacognition, but also these terms remain mysterious (Brown, 1987).

We, thus, believe that research on problem solving, at some point, will have to face squarely the conceptual muddle surrounding introspection, metacognition, and consciousness.

Problem solving research has encountered an impasse. We are fixated on problem solving as search and we are spending most of our (re)search time in problem spaces that do not seem to bring us closer to the ill-defined goal of understanding the general principles underlying problem solving. The work by Newell and Simon (1972) was a real breakthrough that allowed us to investigate the mechanics of well-defined search problems. Contrary to the high hopes in the early days of cognitive science no general problem solving theory emerged. While we can formalize the representations and the steps subjects take in a search problem, we do not understand how they choose between different problem representations, different search algorithms, and different heuristics. An accepted formal theory that can describe how representations and heuristics are attained and adapted is still missing. The same is true for a formal theory of insight. Several people have commented on this state of affairs at a recent workshop that inspired this paper (van Rooij et al., 2011). Two recent papers in this journal also offer an analysis of the current impasse in problem solving research (Batchelder and Alexander, 2012; Ohlsson, 2012). We would like to add to these commentaries that, perhaps, the widespread rejection of introspection among cognitive scientists is one of the factors that is holding back the field.

A computer scientist who had started to do experiments on problem solving told the first author that they wanted to interview their participants about their problem solving strategies for a newly developed set of problems. However, when

they consulted with an experimental psychologist they were advised to focus on collecting hard data, like performance measures, and to avoid verbal reports and introspection. We think that most cognitive scientists will react in the same way if a computer scientist who wants advice on running experiments approaches them. But is this good advice? The argument that we would like to put forward in this opinion paper is the following: cognitive scientists have good reasons to be skeptical about introspection as a research method but self-imposed experimental rigor can get in the way of exploratory research. A systematic introspective study is a good way to get experience with a new experimental paradigm and to explore possible problem solving strategies that subjects might be using in the task. In this way, introspection can help develop new hypotheses and ultimately, perhaps, help overcome the impasse the field is facing. Furthermore, in problem solving research the rejection of introspection is particularly harmful because introspection, in the sense of self-observation, is a crucial component of some problem solving processes. Hence, we will argue that research on problem solving could benefit from being more open-minded about introspection, both as a method and as a cognitive phenomenon.

To avoid misunderstandings right from the start, we are very aware of the big conceptual and methodological problems that introspection raises and readers should not expect that we offer any solutions or easy-to-follow recommendations for running introspective studies. We do not think that there are new arguments or new suggestions in this paper. After all, the controversies surrounding introspection have been raging for a very long time. In fact, our two main points—that (a), introspection can be good for exploratory studies, and (b), introspection is an interesting cognitive process in itself—have often been made (e.g., Dörner, 1979; Ericsson and Simon, 1993, 1998; Ericsson, 2003; Fellows, 1976; Flavell, 1976; Reither, 1979). The aim for this opinion paper is merely to reiterate these two points in the context of problem solving and, hence, to point to one potential block that, if removed, might allow problem solving research to progress.

In the following, we first survey common attitudes towards introspection in cognitive science and try to explain why cognitive scientists are often overly skeptical of introspection. We believe that this skepticism is a major block for problem solving research. In order to remove this block, we will discuss examples of successful exploratory research in cognitive science that used questionable introspective methods. We will argue that, as problem solving research is stuck, it is a good idea to temporarily trade experimental rigor for exploratory introspective methods. Our main argument, however, is that self-observation is a form of metacognition that is an important part of problem solving processes and, hence, needs to be studied as a cognitive process. This is done most easily by explicitly asking subjects to introspect. We will review a selec-

tion of studies that, we think, have already moved in the right direction and we will suggest some ways to proceed.

INTROSPECTION AND COGNITIVE SCIENCE

Cognitive scientists like to compare themselves to particle physicists in that the object under investigation cannot be observed directly. Instead they must conduct cleverly designed experiments that allow them to draw inferences about what is going on inside a black box that they cannot open. These inferences are based on theories that are tested and falsified through rigorous experimental methods.

This self-image is based on a caricature of the history of psychology. For a long time psychology was the business of armchair philosophers. As a consequence, early experimental research in psychology was based on introspection but was a failure due to introspection being unreliable. Behaviorists put psychology on a solid methodological foundation by relying only on observable behavior. In this narrative cognitive scientists save the day by studying the unobservable mind while also being methodologically rigorous.

Although this narrative seems very convincing and can be found in many textbooks, it is not accurate from a historical point of view. The battles between behaviorists and cognitivists get all the attention, but there is also a story to be told about the relationship between cognitive science and the earlier, so-called introspective psychology (Brock, 2013; Costall, 2006; Greenwood, 1999).

As introspection is a difficult term and has different connotations in different contexts, let us first state explicitly which contexts we will *not* be concerned with in the following. We will neither be concerned with the casual introspection of folk psychology nor with analytic introspection, nor phenomenology. The context that is relevant for us here is the experimental work on thinking that students of Külpe established in Würzburg at the beginning of the 20th century. They were the first ones to study thinking experimentally and they did so by means of *systematic experimental introspection*. Briefly, in introspection the subject carries out a task that was set by the experimenter. Directly after completing the task, she explains the course of her thoughts to the experimenter who can also ask specific questions about the subjects' thoughts. The books by Jean and George Mandler offer an excellent introduction to this line of work and also trace its influence on modern cognitive science (Mandler and Mandler, 1964; Mandler, 2007).

As we will also discuss think-aloud methods, let us also, right from the start, distinguish them from introspective methods. Both kinds of methods require participants to give verbal reports and are therefore related and often treated together in the literature. We will follow Ericsson and Simon (1993) and distinguish different types of verbal reports based

on the presumed underlying cognitive processes. When participants think aloud they merely report the thoughts they are having anyway, either in the form of inner speech or in a format that is easily translated into speech. This is achieved by training subjects to utter everything that comes to their mind, however incoherent it is, and by minimizing interaction with the experimenter. Hence, think-aloud methods are often considered to be theoretically and methodologically relatively unproblematic, albeit hard to analyze. In contrast, in introspective methods—where the experimenter questions the subject—there are several additional metacognitive processes: observing thoughts, reflecting on them and ordering them so that they can be explained coherently to the experimenter. These additional metacognitive processes are not understood very well, they may distort and bias the reports, and they may even interfere with the cognitive processes that we want to study. Introspective methods are, therefore, problematic from a theoretical and methodological point of view. With the term *introspection* we lump together the ill-understood and problematic metacognitive processes that are central to introspective methods and that distinguish them cognitively from think-aloud methods.

SURVEY OF COGNITIVE SCIENCE TEXTBOOKS

If one wants to get an overview of common attitudes towards introspection among cognitive scientists, looking at textbooks is a good place to start. We therefore looked at a convenience sample of cognitive science textbooks (see Table 1). We searched the index and the table of contents for the term *introspection* and checked which kind of introspection (folk psychological, analytical, phenomenological, experimental) the

authors discussed and what their attitude was. We then looked more concretely for discussions of the work of the Würzburg school (including Otto Selz) to see whether the tradition that had a positive attitude towards experimental introspection for studying thinking is mentioned at all. As introspective methods are often grouped together with other kinds of verbal reports, such as think-aloud methods or protocol analysis, we also searched for these terms. Furthermore, think-aloud methods can be thought of as an improvement and refinement of the introspective methods that were pioneered by the Würzburg school. Hence, we expected that it is more likely that think-aloud methods, rather than introspective methods, are discussed. The famous body of work on problem solving by Newell and Simon relies heavily on think-aloud protocols¹ and so we also checked whether their work is covered.

Almost all of the books that mentioned introspection as a method dismissed it, mostly based on its unreliability or its uselessness for accessing the unconscious processes that underly cognition. Only four books gave any evidence or serious arguments for these claims. The popular book by Anderson (1990) and the book by Medin and Ross (1992) discuss early works of the Würzburg school. Eysenck and Keane (2000) discuss the problems of introspection but they do not mention the Würzburg school. Instead they discuss the debate between Nisbett and Wilson (1977) and Ericsson and Simon (1980). The only book that gives a discussion of introspection that is historically satisfying from the viewpoint of problem solving is Mayer (1992). This book is, however, unusual in that problem solving takes center stage. The remaining books that mention introspection and dismiss it do not give convincing evidence for the claims that introspection is

Table 1.

In a convenience sample of cognitive science and cognitive psychology textbooks we found that many briefly mention introspection (mostly analytic introspection and seldom the Würzburg tradition) and dismiss it immediately. All of the books mentioned the work of Newell and Simon and most discussed it in great detail. Verbal protocols were not always mentioned together with Newell and Simon and often not discussed at all.

	Introspection	Würzburg	Think Aloud	Newell and Simon
Neisser (1967)				X
Cohen (1977)			X	X
Glass, Holyoak, and Santa (1979)			X	X
Anderson (1990)	X	X		X
Mayer (1992)	X	X	X	X
Medin and Ross (1992)	X	X	X	X
Stillings et al. (1995)	X			X
Green (1996)			X	X
Eysenck and Keane (2000)	X		X	X
Reed (2004)			X	X
Friedenberg and Silverman (2006)	X		X	X
Edelman (2008)	X			X
Bermúdez (2010)				X
Goldstein (2011)	X		X	X

unreliable or useless. We agree with Costall (2006) and Brock (2013) that the case against introspection is historically and factually not as clear as most textbooks make it seem. This suggests that introspection is ignored and rejected in cognitive science not for empirical but for other reasons.²

While all textbooks reject introspection as a method, some of them discuss verbal reports, protocol analysis, or thinking aloud as a methodological improvement over introspection (see Table 1). All of the books in our sample that have been published since the seventies discuss the work of Newell and Simon. They usually discuss *Human Problem Solving* (Newell and Simon, 1972) and give quite a lot of space to means-ends analysis.³

But, surprisingly, not all books discuss how Newell and Simon obtained their insights on human problem solving. They do not discuss Newell's and Simon's research strategy that relied heavily on verbal reports and think-aloud methods. In one chapter of their book they give a meticulous analysis of one protocol of one participant solving one problem, in order to demonstrate how, they think, research in problem solving should proceed. While 9 of the 14 books we looked at mentioned verbal reports only 6 can be said to give a satisfactory account of the role of verbal reports for the work of Newell and Simon (Cohen, 1977; Glass, Holyoak, and Santa, 1979; Goldstein, 2011; Mayer, 1992; Medin and Ross, 1992; Reed, 2004). It is puzzling that only half of the textbooks discuss thinking aloud in detail while all the textbooks talk about problem solving in the Newell and Simon tradition.

INACCESSIBILITY OF COGNITIVE PROCESSES

It is sometimes suggested that a verbal report is just a behavioristic name for introspective data (e.g., Costall, 2006). Despite the success that Newell and Simon had with using verbal reports, some authors might thus feel that methods for the elicitation of verbal reports are not generally reliable, hard to use, or only useful for informal explorations and should therefore not be taught in a textbook. Some authors of textbooks might implicitly side with Nisbett and Wilson (1977) who assessed the reliability of introspection and verbal reports in an influential paper. Their abstract starts with a crushing conclusion: "Evidence is reviewed which suggests that there may be little or no direct introspective access to higher order cognitive processes." They go through a large number of studies and show that the reasons subjects give for their behavior are demonstrably different from the actual causes of their behavior. One explanation for these data is that introspective reports are merely post-hoc rationalizations. Subjects will use their folk psychological theories to explain their own behavior without having privileged access to the actual causes of their behavior.

However, verbal reports are not just a behavioristic name for introspective data (see, e.g., Boring, 1953). Researchers in

problem solving, early on, have tried to distance themselves from introspection as a method but at the same time endorsed thinking aloud as a less problematic and very useful tool (e.g., Duncker, 1945, p. 2). Ericsson and Simon (1980) wrote a reply to Nisbett and Wilson (1977) and tried to set the record straight for cognitive science (see also White, 1980). While the conclusions that Nisbett and Wilson (1977) draw are correct for the studies they looked at, these studies invited post-hoc rationalizations by the subjects. In all of the studies that they reviewed, subjects were questioned *casually* about the reasons for their behavior *after* the experiment. The relevant question, however, is not whether casual introspection is unreliable but under what conditions a systematic and controlled method for eliciting verbal reports can be reliable and useful. Given the progress that was made on human problem solving through the systematic use of verbal protocols, it is no surprise that Ericsson and Simon (1980, p. 247) conclude:

For more than half a century, and as the result of an unjustified extrapolation of a justified challenge to a particular mode of verbal reporting (introspection), the verbal reports of human subjects have been thought suspect as a source of evidence about cognitive processes. In this article we have undertaken to show that verbal reports, elicited with care and interpreted with full understanding of the circumstances under which they were obtained, are a valuable and thoroughly reliable source of information about cognitive processes. It is time to abandon the careless charge of "introspection" as a means for disparaging such data. They describe human behavior that is as readily interpreted as any other human behavior. To omit them when we are carrying the "chain and transit of objective measurement" is only to mark as terra incognita large areas on the map of human cognition that we know perfectly well how to survey.

It is instructive to discuss one example of Nisbett and Wilson (1977) that deals with problem solving (the other examples belong to the realm of social psychology and are less interesting for us here). Maier (1931) hung two cords from the ceiling and told subjects to tie them up. The cords, however, were so located that subjects could not reach both of them at the same time. Subjects could use whatever was available in the room to solve this problem. The solution that Maier wanted his subjects to come up with was to use a weight tied to a cord so that it can be set in motion, like a pendulum. Many subjects did not find this solution even after ten minutes of thinking. In these cases the experimenter would walk across the room and "accidentally" brush one of the cords so that it started swinging. 23 of 61 subjects solved the problem very shortly after having been given this cue (24 solved it before

being given the cue and 14 not at all). Afterwards subjects were asked how they solved the problem. Among these 23 only 7 reported the cue although the remaining subjects were equally influenced by it, as Maier argued convincingly. This is definitely an interesting result. It can be concluded that after having found the solution most subjects did not know about the factors that led them there. In fact, these subjects reported that the solution just appeared to them as a whole.

However, it is clearly a mistake to conclude from these data that subjects do not have *any* introspective access to their cognitive processes during problem solving. We do not know whether a better controlled and concurrent think-aloud instruction would have revealed that subjects noticed the cue and what intermediate steps they went through. Perhaps they simply forgot about the cue after it led them to the solution and so the cue was not available for a retrospective report (Ericsson and Simon, 1980). Also, a third of the subjects did notice the cue and reported it. Even if many of the subjects at no point were aware of the cue it would not follow that verbal reports are useless and unreliable in general. Verbal reports just do not tell us everything we want to know: “A protocol is relatively reliable only for what it positively contains, but not for that which it omits. For even the best-intentioned protocol is only a very scanty record of what actually happens” (Duncker, 1945, p. 11). Hence, the conclusion that verbal reports are generally unreliable and cognitive processes are introspectively inaccessible is much too strong.

LIMITED USE OF VERBAL REPORTS IN COGNITIVE SCIENCE

In our experience, such extreme skepticism about verbal reports as put forward by Nisbett and Wilson (1977) will not be found among people who work on problem solving. However, researchers in other fields, for example decision-making, are likely to hold the view that verbal reports are useless for research in cognitive science. They do not need to rely on verbal reports to the degree that researchers in problem solving have to. Hence, unless the author of a textbook comes from a problem solving background, she is unlikely to see thinking aloud as an important methodology in cognitive science. In problem solving, a subject will often sit and think for quite a while before any overt behavior can be observed. Inferring everything that happened in this period from a few button presses is a very hard task. Even if you add eye tracking and brain imaging it will still be hard to figure out what is going on while the subject is thinking. As human subjects can speak, we can just ask subjects to report their thoughts in order to get more data. Researchers in problem solving will generally agree that think-aloud protocols can, at the very least, be useful for exploratory purposes or applied research (Crandall, Klein, and Hoffman, 2006). But researchers in problem solving are also generally aware of the work of Ericsson and Simon (1993) who argue convincingly that verbal

reports can be more than a source of inspiration. Also, there are some excellent demonstrations of this beyond Newell's and Simon's work (e.g., Kaplan and Simon, 1990; van der Henst et al., 2002; VanLehn, 1991).

However, even if a textbook author acknowledged the historical importance of verbal reports for some of the most influential works in problem solving, she might still be skeptical about the usefulness of verbal reports for future cognitive scientists. Verbal reports are not widely used in research on perception, memory, categorization, decision making, or other successful fields of cognitive science. Even within problem solving they so far have not been that useful in overcoming the impasse that the field is facing. Nobody follows the research strategy that Newell and Simon outlined in their book anymore. Nobody tries to derive detailed step-by-step computer simulations from protocol data anymore (Ohlsson, 2012). As textbooks aren't history books that's a good reason to drop a topic or only mention it in passing. Ohlsson (2012) gives several reasons for why Newell's and Simon's approach never fully entered the cognitive science mainstream. First, different participants behave differently and protocol data are not easily pooled and summarized. Second, and related, collecting verbal reports and performing a protocol analysis differs greatly from the usual methodology of experimental psychology and statistical hypothesis testing. Third, the approach means a lot of hard work for the researcher. And fourth, after it was demonstrated that it could be done, “it became less and less clear what was gained by making it yet again” (Ohlsson, 2012, p. 113). In the same vein, Batchelder and Alexander (2012, p. 69) note that for insight problem solving the data that verbal protocols provide beyond other behavioral measures have not led to any deep theoretical insights.

In summary, most basic research in cognitive science has little use for introspection or thinking aloud. Researchers in problem solving have been more open towards verbal reports but after early successes, progress stalled and the insights-to-effort ratio became very unfavorable. Together with the common historical rejection of introspection this could explain why many cognitive science textbooks do not cover verbal reports, not even in their problem solving chapters. This negative attitude towards introspection, and verbal reports in general, might have stopped researchers from using introspection as an exploratory method or studying introspection as a cognitive process. However, both of these aspects of introspection might help overcome the current impasse in problem solving research, as we will argue below.

INTROSPECTION AND EXPLORATION

Köhler in his foreword to the English translation of Duncker's monograph on problem solving (Duncker, 1945) writes: “It will be objected that any new endeavor which makes little

use of quantitative techniques must necessarily lack the precision of science in a strict sense. Let us not be impressed by this argument. . . . How do the critics plan ever to conquer difficult new ground if the absence of paved roads seems to them sufficient reason for keeping out entirely?” Köhler thus defended the think-aloud methods that Duncker used in his first studies on problem solving. Since the work of Newell and Simon (1972) verbal reports based on think-aloud methods have been used to develop and test detailed models of cognitive processing; and Ericsson and Simon (1980, p. 247) have defended the use of think-aloud methods against the “the careless charge of introspection as a means for disparaging such data.” But even methods that are more problematic than thinking aloud, like classical introspection, can play a role in cognitive science. Once we look at introspection in the context of exploration rather than theory testing we can acknowledge the value of introspection.

INTERACTION BETWEEN EXPERIMENTER AND SUBJECT

A central aspect of systematic experimental introspection as practiced by students of the Würzburg school was that subject and experimenter negotiated the protocol. The experimenter did not just record the verbal reports; she actively tried to understand what the subjects were trying to say and at times questioned them on details in order to get a protocol that was as complete and unambiguous as possible (Ach, 1905; Bühler, 1907). Everyone who has looked at think-aloud protocols knows the feeling of wanting to go back to the experiment in order to ask the subjects what they meant with certain utterances. Such clarifying interactions between subject and experimenter are carefully avoided in modern think-aloud methods. Ericsson and Simon (1993, p. xv) warn us that “to guarantee a close correspondence between the verbal protocol and the actual processes used to perform the task, this urge toward coherence and completeness must be resisted.” Ach (1905, p. 17) thought otherwise, knowing that in his method “the experimenter plays a more prominent part than in any other psychological method” (translated by Titchener, 1909, p. 87). He was well aware of the potential problems of suggestion and confabulation that this approach might give rise to and tried to minimize their impact as much as he could. Bühler (1907, p. 306) explicitly agreed with Ach’s critics “that one will indeed have to look for an objective confirmation of the assertions of introspection.” However, in interaction with a participant, a knowledgeable and careful interviewer may be able to explore phenomena more quickly than it would otherwise be possible.

The work by Stenning and van Lambalgen (2004) is an excellent recent example for research that is questionable by standard methodology but psychologically extremely interesting. They revisited the Wason selection task (Wason, 1968) but systematically engaged in a “Socratic dialogue”

with their subjects. An experimenter who understands the logical and psychological issues involved and knows when to ask subjects to clarify their reasoning is invaluable when you want to get detailed insights into what a subject was thinking. The Wason selection task is a textbook example of how much cognitive scientists were able to find out about the black box by doing clever experiments that only observed easily quantifiable behavior. Still, when we read the excerpts of the protocols provided by Stenning and van Lambalgen we cannot help but notice how much richer than propositional logic subjects’ reasoning actually is. Stenning and van Lambalgen document a large number of non-classical reasoning processes in the protocols that are unaccounted for by current models. Although the Wason selection task was designed to test simple propositional reasoning, the protocols provide evidence for many different reasoning schemata, going way beyond the ones that are discussed in the literature. Furthermore, we can read how the subjects struggle to establish an interpretation of the task. We can read how they try to understand what they are supposed to do and we can see the problems they have choosing between different reasoning schemata. Even if some of the subjects’ reports were rationalizations they would provide us with a much richer set of hypotheses about reasoning in the Wason selection task than usually considered.⁴

INTROSPECTING EXPERIMENTERS

While subject and experimenter were interacting in a social situation in the studies of both the Würzburgers and Stenning and van Lambalgen, which is clearly problematic, they maintained the separation between subject and experimenter and the interaction happened in a controlled environment. Many experimentalists will oppose such methods publicly but it is our guess that, in private, many of them introspect themselves—which is methodologically even more questionable—when exploring the phenomena that they study. Instead of sweeping this exploratory part of our research under the carpet we should be open and methodical about it.

Kahneman (2011, p. 6) describes the exploratory part of his collaboration with Tversky in his recent book: “We quickly adopted a practice that we maintained for many years. Our research was a conversation, in which we invented questions and jointly examined our intuitive answers. Each question was a small experiment, and we carried out many experiments in a single day. . . . We believed—correctly, as it happened—that any intuition that the two of us shared would be shared by many other people as well, and that it would be easy to demonstrate its effects on judgments.” Of course, Kahneman and Tversky then went on to do their famous experiments to check whether other people shared their intuitions, and this is a necessary step. However, if we spent more time exploring potential hypotheses through systematic in-

trospection, we might start off with better hypotheses and ultimately conduct more interesting experiments. The routine that Kahneman describes is serving that purpose.

In the context of problem solving research, someone who understands the theoretical issues involved in problem solving might be able to quickly observe and express patterns in her own problem solving behavior that a naive subject would not notice. You might also find these patterns by meticulously working through verbal protocols or even by careful observation of overt behavior but where does the idea to look for these patterns come from? Task analysis is one answer, introspection is another. We do not claim that new and interesting hypotheses about problem solving cannot be found by other means, but introspection might be a particularly efficient and easy way to develop new hypotheses.

It could be argued that instead of leading to better hypotheses, introspection will merely lead us astray. This is an empirical question. While few cognitive scientists ever talk publicly about the informal introspections they engage in before they run an experiment, our own anecdotal experience is that we can quickly generate interesting hypotheses about how subjects perform a task by doing the task ourselves. In this way we can catch problems in an experiment even before we run it a pilot study and we get new ideas for data analysis. Often we find that there are strategies that we did not think of before. This will be helpful if you want to constrain the task in a way that all subjects follow the same strategy or if you want to sort the subjects into groups that use the same strategy for an analysis. It can turn out that a strategy we thought to be plausible based on a task analysis requires too much working memory and cannot be executed in a pure form. Or we find that we try different strategies and switch back and forth. Often we find that we can do the task but we cannot discern any strategies that we become aware of easily. If we are worried that what we found in introspection may be idiosyncratic we will ask a colleague to do the task and question her about how she did it.

SYSTEMATIC EXPLORATION WITH INTROSPECTIVE METHODS

Why should we not discuss with each other how we think we solved a problem in order to generate plausible hypotheses about the underlying processes? Why should we not try to develop a methodology for systematic introspection that can be used in exploratory studies? Of course, we have to be aware that these are only exploratory studies. But this does not mean we cannot be systematic and methodical about exploration. As there seems to be the general feeling that problem solving research is stuck (Batchelder and Alexander, 2012; Ohlsson, 2012; van Rooij et al., 2011), perhaps, there is a need for more systematic exploration. Luckily, there are many systematic exploratory methods to be found in applied research that we can build on (e.g., Crandall et al., 2006).

More concretely, several authors have recently suggested to collect a large set of standardized insight problems and use them to get a better empirical base in order to understand insight problem solving (Batchelder and Alexander, 2012; Chu and MacGregor, 2011). Having standards in terms of problem difficulty and problem similarity will certainly be useful for designing and comparing studies. However, it is not immediately clear what kind of studies can reveal the cognitive processes involved in insight problem solving. One way forward could be a large-scale and systematic exploratory study. As cognitive scientists we have standard methods for hypothesis testing, but what is our methodology for exploration? An obvious suggestion is to collect a large database of think-aloud protocols for a set of standardized and well-understood insight problems and then work through the protocols in the hope of finding something interesting or being able to systematize problem solutions or search strategies. Ohlsson (2012, p. 112) has described this approach to be like “natural history” where “you collect interesting specimens, dissect them carefully, and report what you find.” However, this approach means a lot of effort for an unclear outcome, especially as it is not even clear how the massive amounts of data will be analyzed and there are no obvious ways to automate this process.⁵

Also, as noted above, problem solving protocols have been collected in the past (though not on a massive scale) and have not led to many deep insights beyond what we know already. Still, we would like to browse through such a database of protocols for inspiration. Collecting such a database would probably require a huge collective effort of the whole field.

Other possibilities are also laborious but perhaps easier to implement. For example, using a large set of insight problems, one could screen for particularly successful problem solvers and interview them about their strategies. Some of them might have interesting hypotheses about how they do it. Some might even have explicit conscious strategies that can serve as a starting point for further investigations.

On a more fine-grained level, we can also follow in the footsteps of Stenning and van Lambalgen (2004) and engage in a Socratic dialogue with participants who are solving insight problems. Think about this as online data analysis by the experimenter. By carefully listening to the subject and asking the right questions, a knowledgeable experimenter can identify interesting phenomena in the moment that they occur. In this way the experimenter can try to clarify what was going on immediately after something interesting happened. In the second edition of their book, Ericsson and Simon (1993, p. xvi) point out that their earlier criticism of retrospective reports does not apply to reports that are collected immediately after a mental event happened. Hence, carefully conducted interviews might be less problematic than previously thought. For exploratory purposes such a method could be

vastly more efficient than transcribing complete protocols and analyzing them offline with no way to later clarify crucial utterances. We think that this is a good reason to try and develop such methods further.

Another possibility is for researchers to solve many insight problems themselves, introspect or think-aloud while doing it, and keeping voice recordings and meticulous notes about what they could observe. Assuming that, with some practice, it is possible to work on a problem and intermittently introspect what you are doing, a knowledgeable researcher might quickly discover interesting conjectures about her own problem solving behavior. If several researchers solved the same problems they could compare their observations and might thereby achieve some limited degree of objectivity.

We fully understand the reservations that cognitive scientists have against such proposals. However, lacking an established and accepted methodology for exploratory studies in cognitive science, we will have to try different methods, compare them to each other, and develop them further. Given the current theoretical block in problem solving research, we think that there is a need for systematic exploration. The tools that are available, think-aloud protocols and psychometric methods, are rigorous but inefficient. At the very least, introspective methods can trade rigor for efficiency. Moreover, instead of outright rejecting the possibility that introspection might be useful as a method to investigate problem solving, perhaps we should investigate when it is helpful and when misleading. In their efforts to develop and understand think-aloud methods Ericsson and Simon (1993) have also documented what we know about introspection. Their framework for interpreting verbal data can also help us interpret introspective data. We are far from understanding the processes involved in introspection completely but we know a lot more about them than the Würzburgers did 100 years ago. “We surely could proceed more safely if we knew precisely in how far we can trust introspection; but how will we find out if we do not try it out?” (Bühler, 1907, p. 306)

INTROSPECTION AND METACOGNITION

We have argued that introspection is useful as an exploratory tool in problem solving research. This is a methodological reason to care about introspection. But we also think that the cognitive processes involved in introspection are key to understanding problem solving behavior. Our own experience as problem solvers suggests to us that conscious self-observation, self-monitoring, self-reflection, and other metacognitive processes related to introspection play a big role in problem solving—especially, when it comes to solving new and hard problems. As we will review below, it has been found in several studies that the instruction to introspect improves subjects’ problem solving performance, sometimes dramati-

cally. Trying to understand the metacognitive processes underlying these improvements is a promising approach for understanding problem solving in general.

REACTIVE AND NON-REACTIVE VERBALIZATION

We have noted above that there are differences between thinking aloud and introspection and that the former seems methodologically less problematic than the latter. Particularly relevant for problem solving research is the *cognitive difference* between introspective and think-aloud methods. This cognitive difference is usually discussed in the context of the question of whether the method that we use to study thinking interferes with the processes of thinking that we want to study. It is not hard to find quotes in the literature that dismiss introspection for this reason. For example, Hull (1920, p. 8) in discussing the work of Fisher (1916) states: “It is difficult to say what influence her constant and elaborate introspections had upon the process, though it is safe to assume that such an amount of irrelevant mental activity was not without its effect.”

As this is a serious concern, Ericsson and Simon (1993) in their defense of think-aloud methods had to convince other researchers that think-aloud methods do not interfere with the normal course of thinking. In other words, they had to show that the verbalizations obtained under think-aloud instructions are not *reactive*. By varying instructions and comparing performance measures with and without verbalization it is possible to study the influence of verbal reports on thinking. Ericsson and Simon (1993) reviewed several such studies and argued convincingly that think-aloud methods are indeed not reactive and the performance with and without verbalization was very comparable in most cases. A recent meta-analysis strongly corroborates this conclusion (Fox, Ericsson, and Best, 2011).

The key for non-reactive verbalization is that subjects only utter the thoughts that they are having anyway, probably in the form of inner speech. However, when subjects explain what they are thinking to the experimenter in a social situation (as in the study of Stenning and van Lambalgen, 2004) or when the experimenter as a subject explains to herself what she is thinking, subjects will have to think about their own thinking. Hence, in introspective methods subjects will attempt to consciously reflect on their thoughts in order to understand, and be able to explain their reasoning. The instruction to introspect triggers metacognitive processes. We imagine that subjects pause from time to time to reflect on what they have just thought and how they have proceeded. They will try to organize their thoughts so that they can be communicated clearly. These reflections may well change the course of thinking. Therein we see the cognitive difference to think-aloud instructions where subjects “merely” report what they are thinking without further reflection. In fact, it

can be demonstrated experimentally that contrary to think-aloud instructions the instruction to explain one's thinking—which is a part of introspective methods—is reactive and changes subjects' performance on problem solving tasks.

POSITIVE EFFECTS OF THE INSTRUCTION TO EXPLAIN

When Ericsson and Simon (1993) reviewed studies that used think-aloud methods and argued that they were not reactive, they also reviewed studies that used other instructions for verbalization. They found that the instruction to explain one's thoughts to the experimenter usually has a positive effect on problem solving performance and a recent meta-analysis confirms this impression (Fox et al., 2011).⁶

Explaining your reasoning or your reasons to behave in a certain way is an important aspect of introspective methods. It is this aspect that is most criticized and which thinking aloud improved on. From a methodological point of view it is questionable to use introspective methods because they interfere with subjects' "normal" thinking (not to speak of the dangers of suggestion, confabulations, rationalizations, and post-hoc folk-psychological explanations). But as a phenomenon in problem solving it is intriguing that performance can be improved by asking subjects to explain their reasoning to the experimenter. Given the big positive effects that introspection can have on problem solving, it is no surprise that it has been suggested that educational psychology could use these results to improve problem solving performance (e.g., Dominowski, 1998; Ericsson and Simon, 1993; Reither, 1979).

For example, the Wason selection task with concrete materials leads to more choices that are consistent with material implication than the abstract version. Even if subjects work on the concrete version first, transfer to the abstract version will usually be poor. However, subjects who have to explain their reasoning to the experimenter show a substantial transfer effect (90%) whereas subjects who are not required to explain their reasoning do not (only 27%) (Berry, 1983). It is quite conceivable that introspecting and explaining made subjects aware of the possibility of transfer. Noticing the relevance of a hint or the possibility for transfer is a crucial factor in successful problem solving (Gick and Holyoak, 1980).

The effect of verbalization on problem solving has also been studied with the Tower of Hanoi problem (Gagné and Smith, 1962). Subjects usually practice with easy versions and are then asked to solve hard versions in a transfer test. Subjects who are asked to give reasons for each move are much better at finding the general principle required to solve the hard versions. In one study, the close to optimal performance of the 14 subjects who gave reasons suggests they all acquired a sufficient, if only partial, understanding of the principle. The other 14 subjects who did not have to give reasons for each move showed a much worse performance that makes it extremely doubtful that they understood the principle. A

hand-waving explanation for the increased performance is that "verbalization somehow forced the [subjects] to think" (Gagné and Smith, 1962, p. 17).

Trying to verbalize one's thoughts presumably helps to organize one's thoughts. This would be consistent with everyday experience that strongly suggests that people often verbalize while they think. It would also be consistent with teachers' recommendations to explain material to someone else or to yourself. Students who explain material to themselves show a better performance on transfer tasks than students who do not. In addition, prompting students to generate explanations improves understanding (Chi, De Leeuw, Chiu, and LaVancher, 1994; Chi, Bassok, Lewis, Reimann, and Glaser, 1989). However, it is not verbalization alone that improves performance. Obviously, the generation of explanations is a more complicated process than mere verbalization as in thinking aloud. Reither (1979) asked participants to self-reflect, rather than to explain, and he also found that performance improves (a summary in English can be found in Dörner, 1979). This suggests that metacognitive processes play a key role for the observed improvements. In some think-aloud protocols that were collected while students learned from worked examples around 40% of the statements are self-monitoring statements. Good problem solvers do not necessarily show relatively more self-monitoring statements than poor problem solvers but they report more comprehension failures, suggesting that they are better at self-monitoring (Chi et al., 1989). Of course, for hard problems poor problem solvers quickly realize that they do not understand anything and hence they also show comprehension failures, so what is the difference? Some of the good problem solvers who spontaneously generated self-explanations systematically used the worked examples to self-monitor their understanding; the poor problem solvers did not do that (Renkl, 1997). Probably, good problem solvers who are systematic about self-monitoring notice *specific* comprehension failures that help them diagnose problems and they have strategies to do something about them.

Later experiments with the Tower of Hanoi problem also demonstrated that verbalizing per se does not improve problem solving but triggering metacognitive processes does. In particular, paying attention to the processing level and becoming aware of what one is doing seem to be important. Berardi-Coletta, Buyer, Dominowski, and Rellinger (1995) asked participants to think about the following questions: "1. How are you deciding which disk to move next? 2. How are you deciding where to move the next disk? 3. How do you know that this is a good move?" These questions can be thought of as asking subjects to explain their reasoning but they can also be considered as an instruction to introspect systematically. Interestingly, participants' performance will improve even if they do not verbalize their thoughts. When they do, however, it can be seen in the protocols that participants who were asked to

think about the processing level focused more on planning, sub-goaling, error monitoring, strategy evaluation, and strategy modification than control groups. Presumably the focus on these metacognitive activities led to the increase in performance (Berardi-Coletta et al., 1995).

METACOGNITIVE PROCESSES IN STRATEGY ACQUISITION

The obvious question to ask now is: can we describe the presumed metacognitive processes underlying the improved performance in more detail? As Ericsson and Simon (1998, p. 182)⁷ put it: “Thinking aloud has now gained acceptance as a central and indispensable method for studying thinking . . . , and it is time to start examining the mechanisms mediating alternative “reactive” modes of verbalization, such as giving verbal descriptions and explanations of one’s thinking ” in order to understand the “transforming power of reflective thought.” Reflective thinking can be triggered by the instruction to introspect and improves problem solving. Hence, the additional metacognitive processes that are involved in introspection are often not “irrelevant mental activity.” They are an important part of the cognitive processing that we want to study in problem solving research (Dörner, 1979; Reither, 1979).

In the Tower of Hanoi problem explaining one’s thoughts is particularly useful in the beginning when subjects have to assemble a strategy but before skilled and automatic processes can be used (Ahlum-Heath and Di Vesta, 1986). Can we describe and model the meta-strategies that successful problem solvers use when they assemble a strategy? Based on the studies by Berardi-Coletta et al. (1995) and Chi et al. (1989) it is reasonable to assume that conscious introspection, as a cognitive process, plays a central role in inventing problem solving strategies when difficult new kinds of problems are encountered. Once strategies have been assembled, they can be automatized and perhaps even be executed unconsciously. Still, self-observation, self-monitoring, and self-reflection—that is, introspective processes—as triggered by the instruction to explain, often seem crucial for improving one’s strategies and for coming up with a sensible strategy to start with. Consciously realizing that one is stuck or that one has had an insight can lead to decisions to change one’s strategy. To do so sensibly, you have to have the ability to monitor and control the execution of your problem-solving algorithms. We imagine this to be a little like debugging or optimizing your code in programming (Sussman, 1973). You systematically test your program and realize that your algorithm gives a wrong answer, is too slow, gets stuck in a local maximum, or is not terminating after the time you are willing to wait. What do you do? You look at the execution trace to identify the problem and then modify the program accordingly. If these speculations about the importance of introspective processes for problem solving were to be considered seriously, how should one study these metacognitive processes?

Anzai and Simon (1979) have analyzed the think-aloud protocol of one subject in the Tower of Hanoi problem, described her different strategies at different time points, and made some concrete suggestions about how the strategies were acquired as a consequence of earlier behavior. They worked in a production system framework. In production systems “re-programming” means acquiring new production rules, so they asked themselves how new production rules could be acquired. They found it necessary to add rudimentary metacognitive processes that monitor success, access execution traces, and generate new rules. Although their system can simulate the transitions between strategies, this cannot, however, be taken as evidence that this is how the subject acquired her strategies. VanLehn (1991, p. 2) noted that models of strategy acquisition are “vastly underdetermined” and that many other models have been suggested since the original work of Anzai and Simon (1979). Ideally, we would like to have more fine-grained, process-level data about the strategy acquisition process which, presumably, is what differentiates good from bad problem solvers. In the language of developmental psychology, we need microgenetic studies (Siegler and Jenkins, 1989). We would like to repeat what Newell and Simon have done for analyzing weak search methods by detailed protocol analysis for strategy acquisition methods.

In this spirit, VanLehn (1991) has tried to find the rule-acquisition events in the protocol that was previously analyzed by Anzai and Simon (1979). He found that the vicinity of the first use of a production rule in the protocol did not give any hints as to how the rule was acquired. VanLehn (1991) argued that the subject approached the problem of strategy discovery in a scientific way. She developed hypotheses and tested them. For example, sometimes she would deliberately ignore established rules to try new ones. VanLehn could identify several rule acquisition events. Some of them were accompanied by clear statements that directly shed light on the process, but most were merely reflected as pauses in the protocol or “visible excitement . . . as if [the subject] knew that she had discovered something general about the puzzle” (VanLehn, 1991, p. 37). As usual, the think-aloud protocols were incomplete. They still suggested several mechanisms and VanLehn could give quite detailed reconstructions of some rule-acquisition events.

The approach of analyzing protocols in order to develop detailed strategy acquisition models has been reasonably successful. In a later paper, VanLehn (1999) fitted his *Cascade* model of strategy acquisition and self-explanation to the think-aloud protocols of Chi et al. (1989). *Cascade* falls into the class of *impasse-repair-reflect* models that have rudimentary metacognitive abilities but are considerably simpler than a full-fledged metacognitive architecture (Cox, 2005; Cox, Oates, and Perlis, 2011; Sun, Zhang, and Mathews, 2006). VanLehn (1999, p. 105) concisely summarizes his work:

The results bring both good news and bad news for Cascade and similar impasse-driven learning models. The good news is that there are indeed learning events, and many of them can be modeled by the simple impasse-repair-reflect cycle. Participants seldom speak coherently about their learning events, but there is evidence for them even in verbal protocols, which are notorious in their incompleteness.

The bad news is that the impasse-repair-reflect cycle is too simple to model all the observed learning, and even the learning-as-self-debugging paradigm might be too simple. Human students appear to be capable of doing everything that good programmers do in developing their programs, except that students are developing their own knowledge. That is, good students can debug their knowledge and can even plan and execute tests that will detect flaws in their knowledge. Moreover, the literature suggests that students can notice opportunities for optimizing their knowledge and perhaps even use their knowledge in a kind of reflective, single-step mode that simultaneously checks the knowledge and executes it.

One way to overcome the impasse in problem solving research could be to collect more process-level data on self-programming and self-debugging. Otherwise it will be very hard to develop computational models of these crucial problem solving processes. Obtaining these data will be a lot harder than collecting think-aloud protocols and creating problem-behavior graphs as Newell and Simon did it. The main reason is that subjects spend a considerable amount of time executing strategies and only comparatively little time on developing new strategies. As the interesting metacognitive episodes are short and rare, think-aloud protocols give only very sparse information about the processes of strategy acquisition. VanLehn's work was successful but he had to comb through hours and hours of protocols to find a very small number of interesting events.

There are two reasons why VanLehn's approach has not been copied very often and the field has not made much progress on finding the mechanisms underlying self-programming. First, self-programming requires self-observation, self-monitoring, and self-reflection, that is, introspection. This will keep many cognitive scientists who are skeptical of introspection from studying these processes. Second, even if someone wants to study these processes, the effort of collecting hours of verbal protocols to get a little bit of data will be immense.

Both blocks can be removed if we embrace introspection and use introspective methods more aggressively. We could use the Socratic methods outlined above as online data analysis. Perhaps, if we asked the subjects immediately after they display "visible excitement" of understanding to report

more details, we could get more data at the time-points that we are most interested in. Alternatively, we could ask expert problem solvers to introspect and explain their strategy-acquisition strategies while they solve many new problems and develop many new strategies. If the introspective methods are reactive, so what? If this increases the frequency of interesting strategy acquisition events in the protocols, we should embrace these methods. In general, we want to make sure that subjects actually engage in the kind of metacognitive processing that we assume to underly successful problem solving. Hence, as one means to that end, we should not be afraid of using reactive methods and explicitly instruct problem solvers to introspect and explain.

CONCLUSION

As research on problem solving appears to be stuck, our suggestion is to relax the strict methodological criteria that many cognitive scientists impose on themselves, at least for a while. We suggest that we should spend more time on exploratory studies using introspection and verbal reports. Getting a better feel for the problem solving phenomena that we want to explain probably has to predate an elegant theory and rigorous experiments. We should, however, not make the mistake of going back to casual and undocumented introspection in single researchers. As a field we will have to be systematic and methodical about introspection if we want to make progress in this way. This will require us to understand the processes underlying introspection and to develop methodologies for reliably collecting introspective data.

Revisiting introspection is not only a methodological recommendation that we want to make. The widespread rejection of introspection has also stopped us from considering its cognitive role in problem solving. However, introspective processes are often crucial for successful problem solving. In particular, strategy acquisition depends on self-observation, self-monitoring, and self-reflection. There is a large literature on metacognition that we can build on, and some of it, especially in educational psychology, even relates to problem solving (Bransford, Sherwood, Vye, and Rieser, 1986; Brown, 1987; Davidson, Deuser, and Sternberg, 1994; Davidson and Sternberg, 1998; Roberts and Erdos, 1993). Metacognition is a much wider term than introspection, but self-observation, self-monitoring, and self-reflection are clearly key aspects of metacognition. As reviewed above, the instruction to introspect can have a big positive effect on problem solving performance. Trying to understand how exactly these improvements come about seems to be a very promising strategy for overcoming the current impasse in problem solving research. Probably the improvements are not due to introspective processes alone but also involve other cognitive and metacognitive processes. The instruction to explain one's thinking is a

trigger for all these processes that in many studies we want to avoid because they change the “natural” course of thinking. But in the study of problem solving, particularly if we want to shed light on how exactly humans invent problem solving strategies, we should be very interested in the “transforming powers of reflective thought” (Ericsson and Simon, 1998). Hence, if we want to study successful problem solving, it will be a good idea to make sure that subjects engage in reflective thought, for example, by instructing them to introspect. Instead of dismissing introspection as a flawed methodology, we should study the role of introspection in problem solving.

ACKNOWLEDGMENTS

The authors would like to thank Felix Wichmann and Sascha Fink for helpful comments.

NOTES

1. As Alan Newell once put it: “As soon as we got the protocols they were fabulously interesting. They caught and just laid out a whole bunch of processes that were going on. . . . My recollection is that I just sort of drew GPS right out of Subject 4 on Problem D1—all the mechanisms that show up in the book, the means-ends analysis, and so on” (McCurdock, 1979, p. 212). We thank Sashank Varma for pointing us to this quote (see also Varma, 2011).
2. Some might conjecture that the main reason is that the work of the Würzburg school is not accessible in English. However, useful summaries have been available for a long time (Boring, 1953; Humphrey, 1963; Mandler and Mandler, 1964).
3. Neisser (1967) is, of course, an exception but he mentions Newell’s and Simon’s early work in passing.
4. For another example see Fellows (1976) who reports that talking to his subjects changed his view on three-term series problems.
5. Titchener (1909, p. 89f) already noted in a footnote to a discussion of the systematic experimental introspection of Ach (1905): “Messer’s paper fills 225 pages of the *Archiv f. d. ges. Psychologie*, and at least half of these are in fine print. There can be no doubt that the method of ‘systematic experimental introspection,’ whatever its advantages, runs to bulk. If it comes into general use, and still more if, as Ach proposes, the conversations between experimenter and observer, the introspective interviews, are taken down by phonograph and stored for future reference, we shall be forced to employ a staff of ‘introspective computers’ to render our materials manageable.”
6. It should be mentioned that there are also cases where problem solving performance deteriorates with verbalization. Thoughts might not be easy to verbalize and trying to verbalize them can interfere with non-verbal thinking, an effect that is called *verbal overshadowing*. Insight problems have been studied most in this context and the data seem to be mixed (Chu and MacGregor, 2011). Still, it has been found that for some insight

problems explaining one’s reasoning decreases the probability of solving the problem (Schooler, Ohlsson, and Brooks, 1993). But the effect is small compared to the large positive effects reviewed here (see also Fox et al., 2011).

7. Thanks to Oke Martensen for referring us to this paper and this quote.

REFERENCES

- Ach, N. (1905). *Über die Willenstätigkeit und das Denken*. Göttingen: Vandenhoeck & Ruprecht.
- Ahlum-Heath, M. E., & Di Vesta, F. J. (1986). The effect of conscious controlled verbalization of a cognitive strategy on transfer in problem solving. *Memory & Cognition*, 14(3), 281–285. <http://dx.doi.org/10.3758/BF03197704>
- Anderson, J. A. (1990). *Cognitive psychology and its implications* (3rd Ed.). New York: W. H. Freeman.
- Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, 86(2), 124–140. <http://dx.doi.org/10.1037/0033-295X.86.2.124>
- Batchelder, W. H., & Alexander, G. E. (2012). Insight problem solving: A critical examination of the possibility of formal theory. *The Journal of Problem Solving*, 5(1), 56–100. <http://dx.doi.org/10.7771/1932-6246.1143>
- Berardi-Coletta, B., Buyer, L. S., Dominowski, R. L., & Relinger, E. R. (1995). Metacognition and problem solving: A process-oriented approach. *Journal of Experimental Psychology: Learning Memory & Cognition*, 21(1), 205–223. <http://dx.doi.org/10.1037/0278-7393.21.1.205>
- Bermúdez, J. L. (2010). *Cognitive science*. Cambridge, UK: Cambridge University Press.
- Berry, D. C. (1983). Metacognitive experience and transfer of logical reasoning. *The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology*, 35(1), 39–49. <http://dx.doi.org/10.1080/14640748308402115>
- Boring, E. G. (1953). A history of introspection. *Psychological Bulletin*, 50(3), 169–189. <http://dx.doi.org/10.1037/h0090793>
- Bransford, J., Sherwood, R., Vye, N., & Rieser, J. (1986). Teaching thinking and problem solving: Research foundations. *American Psychologist*, 41(10), 1078–1089.
- Brock, A. C. (2013). The history of introspection revisited. In J. W. Clegg (Ed.), *Self-Observation in the Social Sciences* (pp. 25–43). New Brunswick, NJ: Transaction.
- Brown, A. (1987). Metacognition, executive control, self-regulation, and other more mysterious processes. In F. E. Weinert & R. H. Kluwe (Eds.), *Metacognition, motivation, and understanding* (Chapter 3). Hillsdale, NJ: Lawrence Erlbaum.
- Bühler, K. (1907). Tatsachen und Probleme zu einer Psychologie der Denkvorgänge. *Über Gedanken. Archiv für die gesamte Psychologie*, 9, 297–365.

- Chi, M. E. G., De Leeuw, N., Chiu, M.-H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, 18, 439–477. http://dx.doi.org/10.1207/s15516709cog1803_3
- Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13, 145–182. http://dx.doi.org/10.1207/s15516709cog1302_1
- Chu, Y., & MacGregor, J. N. (2011). Human performance on insight problem solving: A review. *Journal of Problem Solving*, 3(2), 119–150.
- Cohen, G. (1977). *The psychology of cognition*. New York: Academic Press.
- Costall, A. (2006). Introspectionism and the mythical origins of scientific psychology. *Consciousness and Cognition*, 15(4), 634–654. <http://dx.doi.org/10.1016/j.concog.2006.09.008>
- Cox, M. T. (2005). Metacognition in computation: A selected research review. *Artificial Intelligence*, 169(2), 104–141. <http://dx.doi.org/10.1016/j.artint.2005.10.009>
- Cox, M. T., Oates, T., & Perlis, D. (2011). Toward an integrated metacognitive architecture. In *Advances in Cognitive Systems: Papers from the 2011 AAAI Fall Symposium* (pp. 74–81). Palo Alto, CA: Association for the Advancement of Artificial Intelligence.
- Crandall, B., Klein, G., & Hoffman, R. R. (2006). *Working minds: A practitioner's guide to cognitive task analysis*. Cambridge, MA: MIT Press.
- Davidson, J. E., Deuser, R., & Sternberg, R. J. (1994). The role of metacognition in problem solving. In J. Metcalfe, J. & A. P. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 207–226). Cambridge, MA: MIT Press.
- Davidson, J. E., & Sternberg, R. J. (1998). Smart problem solving: How metacognition helps. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Metacognition in educational theory and practice* (pp. 47–68). Hillsdale, NJ: Lawrence Erlbaum.
- Dominowski, R. L. (1998). Verbalization and problem solving. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Metacognition in educational theory and practice* (pp. 25–46). Hillsdale, NJ: Lawrence Erlbaum.
- Dörner, D. (1979). Self reflection and problem solving. In F. Klix. (Ed.), *Human and artificial intelligence*. Oxford, UK: Elsevier Science.
- Duncker, K. (1945). On problem solving. *Psychological Monographs*, 58(5), 1–113. <http://dx.doi.org/10.1037/h0093599>
- Edelman, S. (2008). *Computing the mind*. Oxford University Press, Oxford.
- Ericsson, K. A. (2003). Valid and non-reactive verbalization of thoughts during performance of tasks. *Journal of Consciousness Studies*, 10(9–10), 1–18.
- Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. *Psychological Review*, 87(3), 215–251. <http://dx.doi.org/10.1037/0033-295X.87.3.215>
- Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data* (Rev. Ed.). Cambridge, MA: MIT Press.
- Ericsson, K. A., & Simon, H. A. (1998). How to study thinking in everyday life: Contrasting think-aloud protocols with descriptions and explanations of thinking. *Mind, Culture, and Activity*, 5(3), 178–186. http://dx.doi.org/10.1207/s15327884mca0503_3
- Eysenck, M. W., & Keane, M. (2000). *Cognitive psychology: A student's handbook*. Hove, UK: Psychology Press.
- Fellows, B. J. (1976). The role of introspection in problem-solving research: A reply to Evans. *British Journal of Psychology*, 67(4), 519–520. <http://dx.doi.org/10.1111/j.2044-8295.1976.tb01542.x>
- Fisher, S. C. (1916). The process of generalizing abstraction; and its product, the general concept. *Psychological Monographs*, 21(2), i–218. <http://dx.doi.org/10.1037/h0093097>
- Flavell, J. H. (1976). Metacognitive aspects of problem solving. In L. B. Resnick (Ed.), *The nature of intelligence* (pp. 231–235). Hillsdale, NJ: Lawrence Erlbaum.
- Fox, M. C., Ericsson, K. A., & Best, R. (2011). Do procedures for verbal reporting of thinking have to be reactive? A meta-analysis and recommendations for best reporting methods. *Psychological Bulletin*, 137(2), 316–344. <http://dx.doi.org/10.1037/a0021663>
- Friedenberg, J., & Silverman, G. (2006). *Cognitive science: An introduction to the science of the mind*. Thousand Oaks, CA: Sage.
- Gagné, R. M., & Smith, E. C. (1962). A study of the effects of verbalization on problem solving. *Journal of Experimental Psychology*, 63(1), 12–18. <http://dx.doi.org/10.1037/h0048703>
- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, 12(3), 306–355. [http://dx.doi.org/10.1016/0010-0285\(80\)90013-4](http://dx.doi.org/10.1016/0010-0285(80)90013-4)
- Glass, A. L., Holyoak, K. J., & Santa, J. L. (1979). *Cognition*. Reading, MA: Addison-Wesley.
- Goldstein, E. B. (2011). *Cognitive psychology: Connecting mind, research and everyday experience* (3rd Intl. Ed.). Belmont, CA: Wadsworth.
- Green, D. W. (Ed). (1996). *Cognitive science: An introduction*. Oxford, UK: Blackwell.
- Greenwood, J. D. (1999). Understanding the “cognitive revolution” in psychology. *Journal of the History of the Behavioral Sciences*, 35(1), 1–22. [http://dx.doi.org/10.1002/\(SICI\)1520-6696\(199924\)35:1<1::AID-JHBS1>3.0.CO;2-4](http://dx.doi.org/10.1002/(SICI)1520-6696(199924)35:1<1::AID-JHBS1>3.0.CO;2-4)
- Hull, C. L. (1920). Quantitative aspects of the evolution of concepts. *Psychological Monographs*, 28(1), 1–86. <http://dx.doi.org/10.1037/h0093130>
- Humphrey, G. (1963). *Thinking: An introduction to its experimental psychology*. New York: Wiley.

- Johnson-Laird, P. M. (2008). *How we reason*. Oxford, UK: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199551330.001.0001>
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Kaplan, C. A., & Simon, H. A. (1990). In search of insight. *Cognitive Psychology*, 22, 374–419. [http://dx.doi.org/10.1016/0010-0285\(90\)90008-R](http://dx.doi.org/10.1016/0010-0285(90)90008-R)
- Maier, N. R. F. (1931). Reasoning in humans II: The solution of a problem and its appearance in consciousness. *Journal of Comparative Psychology*, 12(2), 181–194. <http://dx.doi.org/10.1037/h0071361>
- Mandler, G. (2007). *A history of modern experimental psychology: From James and Wundt to cognitive science*. Cambridge, MA: MIT Press.
- Mandler, J. M., & Mandler, G. (1964). *Thinking: From association to gestalt*. New York: Wiley.
- Mayer, R. E. (1992). *Thinking, problem solving, cognition* (2nd Ed.). New York: W. H. Freeman.
- McCurdock, P. (1979). *Machines who think*. New York: W. H. Freeman.
- Medin, D. L., & Ross, B. H. (1992). *Cognitive psychology*. Fort Worth, TX: Harcourt Brace Jovanovich.
- Neisser, U. (1967). *Cognitive psychology*. New York: Appleton-Century-Crofts.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3), 231–259. <http://dx.doi.org/10.1037/0033-295X.84.3.231>
- Ohlsson, S. (2012). The problems with problem solving: Reflections on the rise, current status, and possible future of a cognitive research paradigm. *Journal of Problem Solving*, 5(1), 101–128.
- Reed, S. K. (2004). *Cognition: Theory and application* (6th Ed.). Belmont, CA: Wadsworth.
- Reither, F. (1979). *Über die Selbstreflexion beim Problemlösen*. Doctoral dissertation, Justus-Liebig-Universität, Gießen.
- Renkl, A. (1997). Learning from worked-out examples: A study on individual differences. *Cognitive Science*, 21(1), 1–29. http://dx.doi.org/10.1207/s15516709cog2101_1
- Roberts, M. J., & Erdos, G. (1993). Strategy selection and metacognition. *Educational Psychology: An International Journal of Experimental Educational Psychology*, 13, 259–266.
- Schooler, J. W., Ohlsson, S., & Brooks, K. (1993). Thoughts beyond words: When language overshadows insight. *Journal of Experimental Psychology: General*, 122(2), 166–183. <http://dx.doi.org/10.1037/0096-3445.122.2.166>
- Siegler, R. S., & Jenkins, E. (1989). *How children discover new strategies*. Hillsdale, NJ: Lawrence Erlbaum.
- Stenning, K., & van Lambalgen, M. (2004). A little logic goes a long way: Basing experiment on semantic theory in the cognitive science of conditional reasoning. *Cognitive Science*, 28, 481–529. http://dx.doi.org/10.1207/s15516709cog2804_1
- Stillings, N. A., Weisler, S. E., Chase, C. H., Feinstein, M. H., Garfield, J. L., & Rissland, E. L. (1995). *Cognitive science: An introduction*. Cambridge, MA: MIT Press.
- Sun, R., Zhang, X., & Mathews, R. (2006). Modeling meta-cognition in a cognitive architecture. *Cognitive Systems Research*, 7, 327–338. <http://dx.doi.org/10.1016/j.cogsys.2005.09.001>
- Sussman, G. J. (1973). *A computational model of skill acquisition*. Cambridge, MA: Artificial Intelligence Laboratory AI TR-297, MIT.
- Titchener, E. B. (1909). *Lectures on the experimental psychology of the thought-processes*. New York: Macmillan. <http://dx.doi.org/10.1037/10877-000>
- van der Henst, J.-B., Yang, Y., & Johnson-Laird, P. M. (2002). Strategies in sentential reasoning. *Cognitive Science*, 26, 425–468. http://dx.doi.org/10.1207/s15516709cog/2604_2
- van Rooij, I., Haxhimusa, Y., Pizlo, Z., & Gottlob, G. (2011). Computer science & problem solving: New foundations (Dagstuhl Seminar 11351). *Dagstuhl Reports*, 1(8), 96–124.
- VanLehn, K. (1991). Rule acquisition events in the discovery of problem-solving strategies. *Cognitive Science*, 15(1), 1–47. http://dx.doi.org/10.1207/s15516709cog1501_1
- VanLehn, K. (1999). Rule-learning events in the acquisition of a complex skill: An evaluation of Cascade. *The Journal of the Learning Sciences*, 8(1), 71–125. http://dx.doi.org/10.1207/s15327809jls0801_3
- Varma, S. (2011). Criteria for the design and evaluation of cognitive architectures. *Cognitive Science*, 35, 1329–1351. <http://dx.doi.org/10.1111/j.1551-6709.2011.01190.x>
- Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, 20(3), 273–281. <http://dx.doi.org/10.1080/14640746808400161>
- White, P. (1980). Limitations on verbal reports of internal events: A refutation of Nisbett and Wilson and of Bem. *Psychological Review*, 87(1), 105–112. <http://dx.doi.org/10.1037/0033-295X.87.1.105>