Patterns, Rules, & Discoveries in Life and in Science

David Klahr

Inspired by the theme of the Festschrift – "From Child to Scientist" – I have divided this chapter into two relatively distinct parts. In the autobiographical part, I describe a series of important events along my path from Child to Scientist. That part is, necessarily, very personal. I reflect on how early events in my life have influenced the way I think and feel about doing science. The second part is about science, in particular, about how the "child as scientist" discovers regularities in the world, encodes and abstracts them, and uses them to make predictions. The broad domain has to do with quantitative development, which happens to be the topic on which I began my career in cognitive development, so that even the second part of this chapter has an autobiographical flavor to it. I will describe an unsolved question about children's thinking, and speculate about how it might be investigated in the future. The question comes from the area of children's early numerical thinking, an area in which some challenging questions remain unanswered.

A child's path to science: from sorting, surveying, satellites, and serendipity at Stanford, to computer simulations.

A Festschrift provides the self-indulgent luxury of reflecting on questions that don't usually come up in the normal day to day plying of one's trade. Questions such as: "What are the forces and experiences in my formative years that profoundly influenced the way I think about my research and the way that I feel about it?" "What satisfies, gratifies, motivates, and excites me, and why?" "What happened to me along the way, and how do those events continue to influence me?" Table 1 lists the answers to these questions in very brief–perhaps cryptic–form; and in this section of the chapter, I will expand and explain each of them.

Category Formation and Parental Approval: Logic and Love

All parents want to find contexts and activities where they can show that they love and support their children. But sometimes it takes a little parental ingenuity to find something that the kids are good at and that the parents also value.

My father had been a pretty good athlete in his youth. When I was growing up, he would occasionally show me a photo, taken around 1920, of his high school football team from P.S. 17, in New York. In the photo, standing directly behind my dad at center, stands quarterback Lou Gehrig; he played football long before he became a baseball icon. My dad clearly valued his early comradeship with someone who went on to become a famous and beloved athlete, and it made his interest in sports very strong. He valued athletics, and he certainly would have encouraged and supported me if I had gotten involved in sports. However, I was definitely not the kind of kid whose natural athleticism was evident from the moment you saw me. As my sister often reminded me during my pre- and early teens, I was chubby, pigeon-toed, near-sighted, timid, and pretty uncoordinated. This certainly did not bode well with respect to any possibility that I could make my dad proud of my performance on the local Little League team (which, as luck would have it, won the world championship in the same year that they formed their first team, and included several kids from my circle of friends!). But my Dad was very ingenious in finding something for me to do that (a) I was good at and (b) he really valued. And that activity had, I believe, a profound impact on my love of science. Let me explain.

Throughout my childhood in Connecticut, my parents owned and operated The Stamford Watch Hospital, a small "ma & pa" jewelry and watch repair business where, according to their newspaper and radio ads, "the sick always recovered". In those days – several years before quartz crystals resonating at 2^{15} hertz replaced the delicate hairspring balance wheels that 'ticked and tocked' –watches were mechanical devices comprised of gears, levers, springs, bearings, and other delicate moving parts that needed to be wound daily, and cleaned, oiled, and adjusted every year or so. After my dad would clean and repair a watch, he would have to let it run for a couple of days to make sure it kept accurate time. But a single winding of the watch only lasted for about 24 hours before the watch would come to a stop, so he had to wind them daily. This schedule meant that, even on Sundays and holidays, he had to go to the store to adjust and wind all the watches that were being checked[1]. On many occasions I would accompany him. There really wasn't much for me to do during the hour or so that it took him to wind, check and adjust all current patients in the Stamford Watch Hospital, but he was very clever. He found a way to keep me occupied, to be of use to him, and to feel good about being useful. Although I didn't realize it at the time, that experience had an enduring impact.

Here's what he would do. He would collect all the excess watch parts on his workbench**:** an accumulated pile of push pins, stems, crowns, springs, hands, bearings, and gears. He would put the pile into an empty cigar box (Dutch Master Blunts, as I recall), seat me at a table in the back of the store, dump the pile on the table, spread it out, hand me a couple of plastic boxes partitioned into a grid of little compartments, and say "sort this stuff". Then he would go to wind and adjust a few dozen "recovering" watches, leaving me to organize the items so that when he needed a specific part, he could find it quickly.

As you might imagine, mapping the multidimensional space that these items occupied into the two-dimensional grid provided by the plastic boxes was a fascinating challenge. How to do it? By material, by function, by size, by shape? Which should be the primary criterion, which the secondary? I used all kinds of different schemes, changing them from one month to the next. When I was done, I would proudly present the sorted and categorized collection to him, describing the latest organizational scheme that I had used. Figure 1 is fairly representative of the result of my efforts.

---

[1] Clearly, this, and many other, aspects of my parents' life as shopkeepers meant that their work was never done. Perhaps this was another kind of early influence on my own career, one that I hadn't recognized until writing this chapter. There is no doubt that academic life includes an unrelenting set of time-urgent demands that blur the line between "working hours" and "non-working hours"-- demands that I, and all of my successful colleagues, seem to have accepted and integrated into our lives.

Figure 1. An array of watch parts, dutifully and creatively sorted.

He always seemed pleased. He would take my sorted assemblage and place it on his workbench, thereby acknowledging that my efforts were, indeed, useful to him. During the next week or so, it was very gratifying for me to come into the store and see my little arrangement at his right hand side on the workbench. My father was not a clinical psychologist … in fact he may have never even heard the term. Moreover, in that era, and particularly in my family, one simply didn't talk about self-esteem, parental approval, self-efficacy, or anything more "psychological" than being in a "good mood" or a "bad mood". But my father was very wise, because he had invented a way to give me an opportunity to discover that the product of my thinking–of my ability to create schemes for classifying and organizing –could be of value and practical use to him … and indirectly to the entire family, because it was of use to the business. I'm convinced that this experience had a deep impact on my psyche. Of course I didn't realize it at the time. And I certainly couldn't have articulated what was happening here … that it was something quite wonderful: cognition, invention, problem solving and precision rewarded by love, approval, and practical utility!

I've been somewhat of a compulsive sorter ever since. In my personal life, it's my tools, my photos, my books, my children's toys: I'm always arranging and rearranging them. Of course, the kinds of cognitive processes that support my little "classification obsession" are essential to the way that we function as scientists. A fundamental part of our work is to categorize, classify, and systematically present our results. Thus, even 60 years later, this early formative experience in my father's shop -- one that engaged both my intellect and affect -- still serves as a source of my enjoyment and satisfaction as a scientist.

Abstract Representations of Reality: Surveys and Maps

The moral of this next story is that when we measure, record, and analyze something in the real world, we create knowledge, but that knowledge is inherently approximate and intentionally abstract. The abstraction process is elegant, and the approximation process is unavoidable. Moreover, participating in both processes can be deeply satisfying, as they were for me when I first experienced them. As in the first example, this resonance is something I did not realize when I first encountered it, but which, as I reflect upon the deeper forces that have kept me on the path of science, seem to have been very important.

Here's the story. When I was in high school and college, I worked after school, and for a couple of summers, as a surveyor's assistant. We did property surveys, ran lines for new roads and sewers, collected data for boundary disputes, surveyed the scenes of traffic accidents, and all the other sorts of things done by survey crews that you see with their tripods, transits, and plumb bobs.

A typical job might be one in which we were hired to make a property map of, for example, the lovely home on a pond depicted in Figure 2a. Imagine yourself in the setting. The grass is green, the birds are singing, the bees are buzzing, the air is hot, the ground is damp, and the pond has a slightly musty smell - a rich setting for the senses. We would arrive with our beat up Willy's Jeep, and I would take the transit out of its box, set it on the tripod, and get the steel measuring tape. My job was to schlep the equipment, hold the "rod" for the guy looking through the transit, cut through brush so as to create a line of sight for the surveyor and his transit, and to do other "grunt" work. But I watched what these guys did, and I was fascinated. From the transit they would read, as accurately as possible from the vernier scale on the circumference of the transit base, the exact angle, to minutes and seconds of arc, of each turn of the transit to the next survey point. Then they would measure the distance from one point in the ground to another (a corner of the lot, an edge of the house or the driveway, etc.) with a 100 ft. steel measuring tape, as one of us held a plumb bob as closely as possible over a survey point in the ground, and we would pull the tape taut to a pre-specified load on a small spring tension measuring device, so that we knew, for example, that we had exactly ten pounds of horizontal force to control for the catenary sag in the tape.

For each measurement, the survey chief would carefully pencil an entry into his battered field book. (No computers in those days!) Only later did I realize that no matter how hard we tried, how careful we were, there was always some error: in reading the vernier scales on the transit, or in locating the plumb bob precisely over a survey point in the ground, or in measuring distances, even with a steel tape and the tension corrections. Of course, I knew that the stuff we did in science lab in school was always full of errors, but I thought of error as a kind of "mistake" rather than an inherent aspect of measurement. I learned an important lesson early, but surely never explicitly articulated in my surveying days, that error is unavoidable in science.

But there was a more important lesson, and it was more abstract, and, for me, more profound. For once we returned to the office, I would watch as the information from the field book was transformed into a map, with the help of straight edge and compass. The challenge was to start at a specific point on the paper, draw straight lines and intersecting angles at scale, such that they would correspond to the "real world" angular and linear measurements. The "holy grail" in this endeavor was to get the end point of the final line on the paper to end precisely on top of the start point of the first line. This "closing the survey" resulted in a lot of satisfaction and pride among the survey crew.

As I observed this process, I was fascinated with the way in which all of our efforts in the field, in the real world, with stumps and bumps, rocks and buildings, and briars and mud, would be transcribed from the field books into maps of the kind shown in Figure 2a, where the house and pond had been transformed into a symbolic abstraction. Much was lost, but *what was essential for the purposes of the survey had been retained*. Moreover, *some knowledge existed in the abstractions that did not exist in the real world.* The distances, angles, elevations, and contour lines, all culminated in a succinct simplification that revealed new relations among the elements.

Figure 2.
a. From the real world to the abstracted and quantified world of surveying.
b. From the real world of children to a spreadsheet and ANOVA results.

Isn't that what do we do as psychologists? We might be studying scientific reasoning, or problem solving, language acquisition, or number concepts; but in all cases, we extract, from the richness of each individual case, only what is of interest to us, and we leave the rest behind. In the kind of work with which I am most familiar, the primary "yield" from many hours of data collection with many children is a spreadsheet with columns for the various conditions and measurements and rows for the children. That is, each child's response to our challenges becomes a row in a spreadsheet. That's all that's left. That's all we want to examine. We have retained what's essential for our purposes and discarded the rest: the children's voices, smiles, cute behaviors, funny but irrelevant comments, and so on. From 50 children to 50 rows in a spreadsheet. And then we abstract again. We take the data, we pour it into our statistics package, and we aggregate and simplify even further in order to tell our story. We present effects, contrasts, and *d*' values. By selective simplification, we have created a new entity, a new kind of knowledge, that did not exist until we did those transformations. The point of this lesson is summarized in Figure 2b. That is, I see a direct analogy between the translation from the physical world to the surveyor's field notes to the final map on one hand, and the real children, our data sheets, and our extracted statistical models on the other.

Why do I find this so interesting, challenging, and satisfying? Well, my exercise in self-analysis is claiming that I was imprinted at a tender age on these aspects and features of surveying because the job came with two powerful affective components. First, it had high prestige amongst my nerdy friends (in other words, all of my high-school friends) because I had been chosen for the position–ahead of my classmates –on the basis of my physics teacher's recommendation as technically competent and reliable. Second, it had high status, even more broadly, because surveying was associated with "macho" construction jobs, with being outdoors, and with working under severe, and occasionally somewhat dangerous, conditions. So the affective aspect was tremendously fulfilling, and the intellectual part is strongly associated with what we do as researchers. That is, the process of doing research is just like what used to happen in my surveying days when we would go from the survey in the damp, muddy, buggy, field to the map or blueprint based on the survey. I'm convinced that my early affective and cognitive experiences as a surveyor's assistant gave me a deeply embedded, although unarticulated, understanding of and attachment to both the elegance and limitations of the research process.

Knowledge Driven Search Trumps Trial and Error

In graduate school, I began to learn about formal models of problem solving and decision making, and about the profound difference in efficiency of "knowledge driven" search over "trial and error" search. I also discovered that I had already had a personal experience in which I had seen this contrast in action, a personal experience that, when I eventually encountered a formal description of it, really rang a bell.

My first job after graduating from college was as a computer programmer working for Wolf Research and Development Company, a very small (~ 10 employees) company in Boston that had several Air Force contracts involving computer programming. My first assignment was to work on what we then called "an adaptive program", but which today would look like some pretty simple machine learning work. That was the sort of thing that had attracted me to the job, because my senior thesis at MIT involved a primitive bit of artificial intelligence – writing a program that learned how to play the game "NIM" by watching an expert play it[2]. However, Wolf also did a lot of "bread and butter" work that was mainly taking data in one format and converting it to another, for example, taking readings from a radar set based on azimuth, elevation, and distance from the radar site and converting it to latitude and longitude. The tasks were pretty straightforward conceptually; but in those days of millisecond machines with only two thousand words of memory, even these mundane tasks took a lot of ingenuity. After I had been at Wolf for a year or so, they landed a big contract with the North American Aerospace Defense Command (NORAD) in Colorado Springs, Colorado, and being young, single, and eager to travel, I jumped at the chance to move west to work on the project.

What did NORAD do? Well, as anyone over 50 or so will recall, these were very serious and crazy times. We were engaged in a "Cold War" with the USSR, and the fundamental military strategy was called Mutually Assured Destruction, or MAD. And mad it was. The basic idea was for each side to guarantee that if one side attacked, the other side would immediately counter attack. Each side knew that they could not intercept the other's nuclear-armed intercontinental ballistic missiles, but they also knew that they could launch enough missiles of their own to destroy the initial attacker, even as they were being destroyed. So nobody wins, and everybody loses. NORAD played a key role in this astoundingly insane zeitgeist, because its job was to determine whether or not anything coming over the horizon was a missile. This decision might not seem to have been much of a challenge, because the United States had enormous radars –

---

[2] When you tire of reading this chapter, try this:
http://www.archimedes-lab.org/game_nim/nim.html#

approximately the size of a football field tipped on its side –sitting in Alaska, Turkey, and England, pointed toward Russia, scanning the horizon.

However, there was a problem, because even in the early 60's a lot of objects -- ranging from exploded rocket boosters [3] to nuts and bolts -- were coming over the horizon every hour, and they were all harmless. Even in the early sixties, there were many objects in near earth orbit. So the big radars peering over the horizon were seeing a lot of moving objects and sending the signals of their tracks to NORAD, at which point our computers would try to determine whether any of these were in a ballistic trajectory–indicating that the Russians had launched their missiles–or in an orbital trajectory, indicative of harmless pieces of metal circling the earth. These computations had to be completed quickly, because it only takes about 15 minutes for a nuclear-armed ICBM to get from launch to target. They also had to be done correctly, because a false negative meant the end of the New York or Washington or our building in Colorado Springs![4] A false positive meant the end of civilization.

The basic computational problem was to match the "track" of the sighted object to either a ballistic or an orbital trajectory. For a single object, this would not have been much of a challenge, even with the existing computational power; but, as I noted above, there were many objects, and thus many tracks to compute … long before the days of parallel computers. Of course, we did not have much computational power … certainly not by today's standards. NORAD's state-of-the-art computer was the Philco 2000: 32K memory, a 1M disk, and 22K multiplications per second. (For the non-technical reader, think of it this way. Your cell phone has about two thousand times as much memory as the computer that was at the heart of the defense system of the "free world" in the 60s.)

The programming teams tried various clever ways to do this discrimination as efficiently as possible. Of course it was all in assembler code, so it was very labor intensive. And then someone had a brilliant idea. So brilliant, and so obvious, that it made a deep impression on me. Here's the idea: *Instead of treating each observation as something totally unknown, make use of what you already know*.

You know that object X is in orbit, and that means you can predict exactly where and when it should come over the horizon in about 90 minutes. So rather than treat each sighting as if you know nothing, once you know what you are looking at on the horizon now, just revise its orbital path a bit, and predict where it should show up the next time around–and you have plenty of time to do it. If, when you look at the first few blips that you think are object X, and those the blips fit the prediction, then you are done with that guy–and all the data associated with that sighting –for another 90 minutes. You just have to make a slight revision to the known orbit. If it’s not there, then it blew up or disintegrated on its last trip around. And that leaves you lots of computational power to focus on the remaining unexpected blips on the horizon. Simple and elegant: *Knowledge trumps brute force computation*. Theory guided search is the way to go!!

I wish I'd thought of that, but I didn't. However, I never forgot the lesson. Always ask yourself, "what do I already know?", before starting a complicated search. Or to put it in terms that Kevin Dunbar, Anne Fay, Chris Schunn and I used in our work on scientific reasoning

[3] Today there are an estimated 20,000 objects at least as large as an apple, and perhaps half a million smaller objects, in near earth orbit. In fact, Vanguard I, launched in 1958, is still in earth orbit. These objects pose an ever increasing danger to space missions.

[4] When I worked at NORAD, it had not yet moved into the "hard site" hundreds of feet underground in Cheyenne Mountain. Our building was called a "soft site".

(Klahr, Fay & Dunbar, 1993; Schunn & Klahr, 1992, 1996), your location in the hypothesis space should guide your search in the experiment space. Little did I know at the time that my experience of peering into "real" space would influence my research in "cognitive spaces". Even today, with all of the incredible computing power available to us, the big advances in computer science come from ingenious formulations of problems, rather than from brute force computation.

Serendipity at Stanford

So much for introspections on early influences. But while I have been focusing on the ways in which specific aspects of my varied experiences have contributed to the attraction and satisfaction of my career as a scientist, I have yet to explain how, given my engineering and programming background, I became a particular kind of scientist - one with an interest in cognitive and developmental psychology. That requires one more personal anecdote, one that was truly transformative, and entirely serendipitous, for it redirected me from one kind of scientific career to another.

The first step on the path to that event was not particularly unusual, so I won't describe it in any detail. It took place in the fall of 1962 when I left the lovely town of Colorado Springs, nestled at the foot of Pikes Peak, and drove to smoky Pittsburgh in my hot little TR-3 sports car, to enter a Ph.D. program in Organizational Behavior in the Graduate School of Industrial Administration (GSIA) at Carnegie Tech (now called the Tepper School of Business at Carnegie Mellon University). I had been attracted to that program because Herb Simon and Allan Newell were at Carnegie Tech as central players in what became called "the cognitive revolution" in the late 50's and early 60's and thus GSIA seemed an ideal place to pursue my long standing interest in doing intelligent things with computers.

After a couple of years of courses in Organization Theory, Economics, and Management Decision Making in a Ph.D. program that Newell and Simon called "Systems and Communication Sciences" (the precursor to what became Carnegie Mellon's School of Computer Science), I had just begun to formulate my dissertation topic on using multidimensional scaling techniques (Kruskal, 1963) to characterize the decision making process of college admissions officers (Klahr, 1969b). But I was still doing background reading and not fully engaged in the work[5]. Along the way, I had learned how to program in one of the then-novel "list programming languages" called IPL-V[6].

Thus, in the Spring of 1965, when I was about half way through my graduate program in GSIA (now Tepper), I happened to be schmoozing with one of the GSIA faculty, Walter Reitiman[7]. I asked him what his summer plans were, and he told me he was going to a 6-week summer conference at Stanford. "Sounds nice," I said. "Want to come?", he asked. "I could use a teaching assistant on how to construct cognitive models in IPL-V." It didn't take a lot of thought

[5] However, I was sufficiently interested in multidimensional scaling to publish a paper on the topic that became one of my most widely cited, even though I never did another psychometric paper (Klahr, 1969a).

[6] This was Carnegie Tech's competitor with MIT's LISP. Although IPL preceded LISP by a couple of years, LISP went on to completely dominate AI programming. Nevertheless, IPL was the language in which many of the landmark programs in AI (EPAM, the Logic Theorist, and the early Chess programs) were created.

[7] Reitman was a true innovator who challenged the seriality of the Newell & Simon approach to cognition by proposing a radically different computational architecture that he called "Argus", inventing, in effect, connectionist computational concepts 20 years before the beginning of PDP modeling (Reitman, 1964, 1965). He was also the founding editor of the journal *Cognitive Psychology* in 1970.

before I agreed. The idea of 6 weeks at Stanford sure sounded nicer than another hot summer in smoky and sooty Pittsburgh, so off I went.

The Conference on Learning and the Educational Process, sponsored by the Social Science Research Council, was decades ahead of its time. Its goal was "to stimulate the thought of any person seriously interested in research approaches to the problems in education" (Krumboltz, 1965, p *ix*). From my perspective ,it more than achieved its goals, because it certainly stimulated my thought, and I was not even interested in "research approaches to the problems in education" at the time! To be honest, I was just looking for a pleasant summer in the Bay area. Suddenly I was thrown into an intense, highly interactive, richly debated conversation with many of the giants (or giants to be) in the field: Robert Gagné, Richard Atkinson, Lee Cronbach, Daniel Berlyne, Jerry Kagan, John Carroll, Bob Glaser, David Premack, Hiroshi Azuma, John MacNamara, Richard Snow[8], among others. Not that I knew they were giants–remember, I was coming from a background, first in engineering, then organizational behavior, with no connection whatsoever to what we now call "the education sciences" – but I was certainly dazzled by the clarity of their thought, the richness of the problems they were talking about, and the importance of the challenges they were addressing. I have remained interested in issues of cognition and instruction ever since (Klahr, 1976, Carver & Klahr, 2001). In fact, I view the Stanford conference as the intellectual precursor of the pre-doctoral training program in the education sciences that my colleagues and I created at Carnegie Mellon half a dozen years ago[9]. But the really profound influence of the Stanford conference was not that it stimulated my interest in educational research, but rather that it turned me toward a career of research on cognitive development.

Here is how it happened. Two or three days into the conference, I initiated a conversation after dinner with a cheerful young Scotsman (with a thick brogue) who had recently completed his Ph.D. in education from the University Warwick in England. His name was Iain Wallace[10], and he opened the conversation with the kind of thing that one does at such conferences,

> “What do you do?” he asked.

I replied, in what I regretfully admit was probably a cocksure tone, “Oh, I write complex computer models of thinking and problem solving,” and then I launched into an extensive discourse on all the wonderful things going on at Carnegie Tech (and they really were quite wonderful, mind you.) I went on for about an hour, at which point I dimly recalled that the conventions of social discourse suggest that, in this sort of situation, you should ask other people what <u>they</u> do.

> Klahr: “Oh, and what do you do?”
> Wallace: “Well, I’m a Piagetian,”
> Klahr: “What’s that?”
> Wallace: “What do you mean, ‘What’s that?’”
> Klahr: “What’s a Piagetian?”
> Wallace: “Oh, someone who studies Piaget,”
> Klahr: “Who? Who’s Piaget?" Isn't there a watch brand called "Piaget"?

---

[8] I estimate that the correlation between a reader's age and the number of recognized names on this list is > .9.

[9] http://www.cmu.edu/pier/

[10] Actually, John Gilbert Wallace, but "Iain" to friends and family.

They say "ignorance is bliss", but such unadulterated ignorance is rare, and I sure had a pure form of it with respect to Piaget and cognitive development at the start of the summer of 1965!

Undaunted, and quite eloquently, Iain began. He told me who Piaget was, and the problems that he was addressing. I learned that Piaget was interested in cognitive structures. I learned about his ingenuous empirical studies with young children (often his own children!). As I listened to Iain's lucid articulation of the fascinating set of phenomena, questions, procedures, and proposed solutions comprising Piaget's "genetic epistemology", I learned that Piaget had formulated his stage theory in the context of a kind of modified algebraic representation, and that his fundamental interest was in cognitive change and dynamic processes. It occurred to me that perhaps Piaget was using an inadequate formalism with which to cast his theory. I said to Iain, "But algebra is the wrong language, because it's static, and computational models are expressed in a dynamic language. Wouldn't it be interesting to try to formulate computational models of the kind of phenomena that Piaget studies: number conservation, class inclusion, and transitivity? Perhaps you and I could collaborate on a project in which we applied 'the Carnegie Tech approach' to problems of cognitive development."

He agreed, and during the remaining weeks of the conference, we began to formulate plans for finding a way for the two of us to collaborate. For me, this was not easy, because I was still on a career track headed toward a faculty position in the Decision Making territory of the business school world, which was not exactly a hotbed of interest, support, or activity in computational models of number conservation! So the challenge for me was to find some funding agencies that would support this new passion of mine for which I had no track record. This took a while, and in the meantime, the inertia of my Ph.D. training kept me on the B-school track. I completed my Ph.D. and took my first position as an Assistant Professor at the University of Chicago, where I taught some courses in Organizational Behavior in the Business School and others on the newly emerging field of "Artificial Intelligence" in the Math Department.[11]

For several years, I tried to get funding from different agencies that would allow me to shift from Organizational Behavior to start collaborating with Iain Wallace on cognitive development research. After several disappointments, we got lucky. In 1968, Wallace secured funds from the British Social Science Research Council for me to spend a semester with him at the newly founded University of Stirling, and I managed to get a Fulbright Teaching Fellowship to teach the following semester at the newly formed London School of Business, which solved the proximity problem[12].

At this point, a slight digression on organizational climate is in order. It was not coincidental that both of the institutions where I spent my first full year collaborating with Wallace were new organizations, and quite open to faculty with non-traditional research interests. Perhaps the message for young faculty is that if you want to shift fields a little bit, look for an innovative institutional context and be willing to take some risks. Also, be prepared to ignore your more conservative colleagues. I know that when, as an assistant professor in my

[11] In the early days of Cognitive Science, universities did not know where to put such topics: some chose Psychology Departments, some chose Electrical Engineering (after all, computers had power supplies, and transistors, didn't they?) and some chose Mathematics Departments.

[12] Of course Stirling, Scotland, and London, England are not exactly in one another's back yards, and there was no internet yet, but for us, this was close enough to continue our projects via several long weekends of collaboration in London or Stirling.

second year at the University of Chicago Business School, I told colleagues that I was planning to take a year off and go to Scotland and England to do research on children's thinking, a typical response was, "That's the stupidest thing I can imagine. You haven't even been reviewed for reappointment yet and you're already leaving, you're already changing fields, you don't have a track record, and you've only published a little bit in your own field. What's the matter with you? Are you nuts?"

Perhaps. But it didn't matter, because while I was still in Scotland, I got a letter from Richard Cyert, then the Dean at GSIA (who later became one of CMU's most influential and respected presidents). Dick, who had been briefed by Herb Simon[13] on how to phrase the job offer, said, "We think that cognitive psychology knows enough now that we could start to engineer better education in business school, so we would like you to come back and be GSIA's 'learning engineer'. Your challenge would be to take what we know about learning and forgetting and memory and problem-solving and get the faculty who teach accounting and economics and marketing to use the results of this emerging discipline of cognitive science to improve their courses." And I replied, "That's very interesting, but I really want to do developmental psychology also, so here's the deal. Suppose I make the 'learning engineering' work my teaching load, while my research focus would be on cognitive development, and I would have a joint appointment in Psychology?" And they said, "Okay. Come back." That explains how I became a faculty member, and in fact, a Department Head for 10 years, in a world class Department of Psychology without ever getting either an undergraduate or a graduate degree in Psychology!

Lest this sound just a little too smug, I need to acknowledge the incredible luck that seemed to embrace me at each fork in my meandering path. An important manifestation of that luck is what Herb Simon called "a secret weapon." He once told me that in order to break into a well-established field from outside you had to have a secret weapon. I was particularly lucky because I had two secret weapons. One was computer modeling. As I noted above, I happened to stumble into one of the few places in the world that was beginning to develop and exploit computer languages for formulating complex theories of cognition, and acquiring that skill several years before it became widely disseminated certainly gave me what economists call a "comparative advantage". My other secret weapon was Iain Wallace: a colleague well trained in cognitive development, and a tremendously creative, energetic, and original thinker.

For the next 10 years or so, Iain and I continued our collaboration in various places: in Scotland, in Pittsburgh, and in Australia (where Wallace eventually moved into an educational administration track). We created production system models for children at different stages of class inclusion, quantification, conservation, and transitivity. Eventually our interests and careers took us in different but similar directions. Wallace became a Dean of Education, and then a Provost at a couple of Australian universities. I moved from a focus on cognitive development to more of a problem solving and scientific reasoning focus, expanding to adults and to educational issues.

I feel very fortunate to have had these two secret weapons. I am convinced that they had more to do with my rapid and successful entry into the field of cognitive development than any extraordinary intellectual skills on my part. There is no false modestly here, because I certainly

[13] Simon's substantial influence on my life occurred at this, and several other crucial points, all of which are described elsewhere (Klahr, 2004).

am aware that I have the requisite "smarts" and energy to have had a reasonably productive career as an academic in one or another modestly technical areas of cognitive science or management science. But the comparative advantage of entering psychology just as the cognitive revolution was gathering its full momentum and of having mastered the requisite computational skills to cast developmental theory in the form of computational models (cf, Klahr & Wallace, 1970a, b) gave my work a kind of instant recognition, that was, in my opinion, well beyond any sort of extraordinary effort or creativity on my part. I was both pleased and confused by all of this. In fact, when I received an invitation to give a talk at the Minnesota Symposium on Child Development in 1971, I did not have a clue about who the people were who invited me (Anne and Herb Pick, two of the most prominent developmentalists of that era); and when I attended the symposium, I spent a couple of somewhat awkward days there, because I did not know a single person, not having come through the normal professional socialization process of a developmental psychology graduate student. However, the appearance of my paper in a very high-prestige symposium volume gave our work even more legitimacy and influence (Klahr, 1973).

I think that these shaping forces and contexts provide at least a partial explanation for why I did what I did and why I do what I do, although I certainly did not view them at the time in the way I have just described them. Only time puts these significant influences into perspective.

Children’s path to number conservation: series completion, subitizing, and statistical learning

In this part of the chapter, I make a transition from autobiography to a scientific question. I will focus on a fascinating puzzle about “number conservation” -- a topic in cognitive development that used to be studied intensively – by several of the authors of other chapters in this volume, in fact –and that was ultimately abandoned by all of us as we moved on to other topics, but without having answered one of its most challenging questions. The puzzle is "how do children acquire empirical evidence about number conservation?" This discussion has three parts. First, I describe the problem. Next I summarize a theoretical account of how children acquire the knowledge elements that enable them to understand number conservation. This account will draw heavily on a theoretical paper that I wrote 25 years ago (Klahr, 1984), a paper that made some claims that, at the time, could not be tested by either the empirical tools or the theoretical models available at the time. Finally, I will suggest that recent advances in our research methodologies and theories make it possible to return to that topic in order to really understand it.

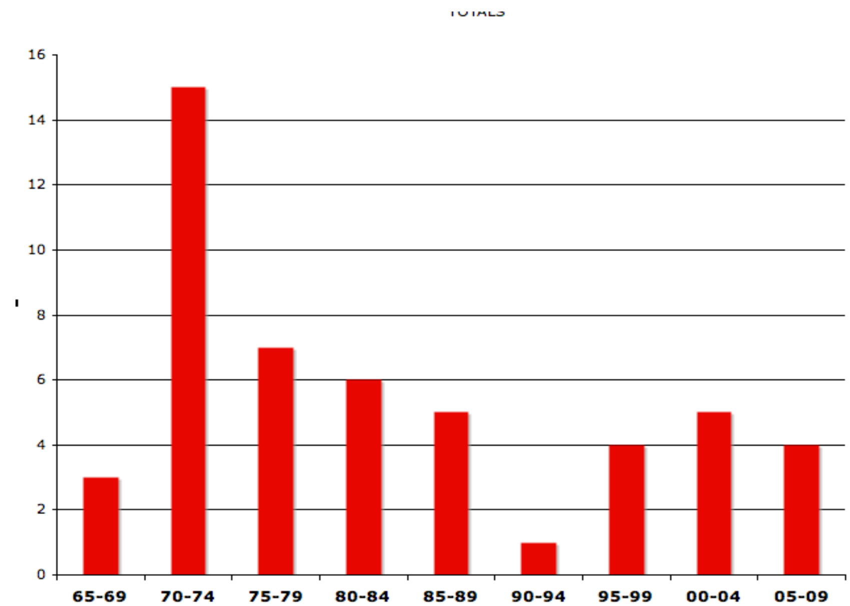

Figure 3. Number of papers published in *Child Development, Cognitive Psychology, Cognitive Science, Cognition, Cognition & Instruction, Cognitive Development, Developmental Psychology, Developmental Science, J. Exp Child. Psych, J of Cognition & Development, J. of Genetic Psych, Psych Bull, and Psych Review, w*ith the words Acquisition of conservation, Acquisition of number conservation, Conception of number, Conservation acquisition, Conservation learning, Conservation of discontinuous quantity, Conservation of number, Conservation of quantity, Conservation skills, Development of number, Inducing conservation, Number concept, Number concepts, Number conservation, Number development, Number invariance, One-one-correspondence, Quantitative invariance, Quantity conservation, Reasoning about number, Subitizing, Training conservation, or Transfer of conservation in their titles or abstracts during each 5 year period from 1965 to 2005.

If you have not followed this field since your undergraduate days, you might think this topic kind of "old fashioned", that which we know pretty much all there is to know. As shown in Figure 3, there was a substantial amount of research on the general topic of number conservation in the early 70's, but it dropped precipitously through the early 90's. However, recent years have seen a resurgence of research on the development of quantitative concepts, with number conservation at the core of that interest. As I will argue below, much of this activity has been stimulated by the theoretical and methodological advances in our field.

What are the Knowledge Elements that Comprise Number Conservation?

If a set contains a certain number of discrete items, and if they undergo transformations such as *spreading*, *rotating*, *compressing*, or *transposing*, then the number of items does not change. That is, the types of physical actions just listed are all invariant with respect to number (*aka* "number conserving transformations"). However, if you *remove, eat, add, subtract, or vaporize* one or more items, then the number of items in the set does change. The second group of actions are NOT invariant with respect to number (*aka* "number changing transformations"). This distinction is so obvious to any normal adult that it seems a bit pedantic to even state it. However, it is not a piece of knowledge that young children have. Even since Piaget's time, tens of thousands of pre-school children around the world have sat across tables from developmental psychologists and then (a) been presented with arrays containing a small number of objects, (b) been asked to quantify the amounts in each collection, and/or asked to determine their relative numerosity, (c) observed as one or both arrays was subjected to one or more of the types of transformations just listed, and finally (d) been asked to make a statement about the relation between the initial number of set elements and the final number, or between the transformed set and the untransformed set.

The specific attributes of this canonical procedure vary along several dimensions: one set or two; the number of objects; whether one set had more, less, or the same number as the other set; the type of transformation; the spatial layout of the arrays; the heterogeneity or homogeneity of the objects; the semantic relations between one set and another (e.g., all pennies in both sets, or eggs in one set and egg cups in the other); and so on. For example, Figure 4 shows a set of typical variants of the type studied by many conservation researchers to explore the relation between physical arrangement, type of transformation, and set size. In Task 1, the two arrays start the same with respect to number, length, and density. Then array B is compressed, reducing its length and increasing its density, but not changing number. Finally, the child is asked to judge the relative numerosity of sets A and B'. In Task 3, array A starts out with less numerosity, length and density than array B; B is then compressed so that it is more dense, and equal in length to A. Again, the child is asked to judge the relative numerosity of the two sets. The wording of the final question is yet another variable in this type of research. The experimenter could ask a vague question, such as "which collection is bigger?", to which the child might respond on the basis of length rather than number; or the question might be more focused on number–albeit still a bit vague- such as "are there more here, or here?"; or the question could be even more explicit about the fact that it is number, rather than length or density that is the focal dimension, as in "which set has more items?" or "who has more cookies to eat?" But this study is just one of many, and the size of the experiment space (Klahr, 2000) here has sustained a small industry of developmentalists, and produced a vast empirical base about the conditions under which children appear to understand the difference between number-preserving transformations and number-changing transformations.

From the perspective of any typical adult, all of these variants might seem irrelevant. All I need to tell you about one of these experimental manipulations is the initial relation and the type of transformation. If sets A and B start with an equal number of objects, and I tell you that I moved the items in set B closer together, you know that A and B have remained numerically equivalent. If I tell you that I ate one of the cookies in set A, you know that set A is now smaller than set B. You need not look at the final collections to determine their relative size. Why? Because you know that compression has no effect on number and eating does. To reiterate this crucial point: a child who fully understands conservation need only know the initial state and transformation type. There is no need to encode the final arrays, because transformational types uniquely determine the outcome[14].

[14] Ah yes, the reader will interject, "but what about ambiguous cases, such as where one set starts larger than the other and the transformation is quantity changing? Without exact numerical information about the *size* of the transformation, the outcome is ambiguous." Granted, but not crucial to the argument here.

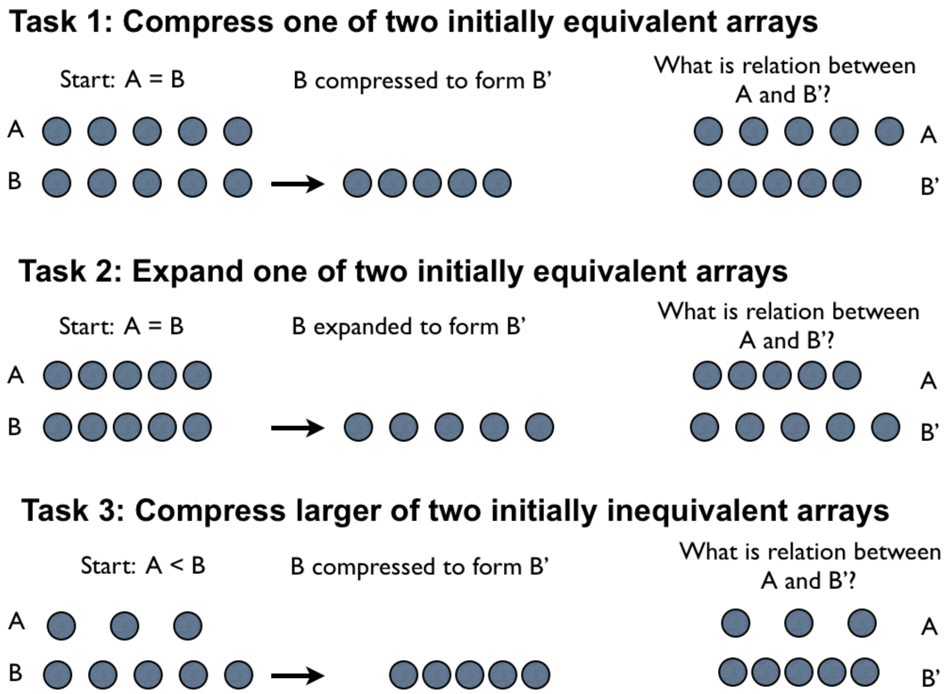

Figure 4. Typical conservation tasks

In her path-breaking study of children's conservation knowledge, Gelman (1972) demonstrated that 4 year olds were surprised when what they expected to be quantity-preserving transformations – Gelman called them "number irrelevant transformations" – yielded apparent changes in quantity (because the experimenter was doing "magic tricks" with the objects).

> "As the magic experiments demonstrate, when children reason about numerosity they recognize the existence of a large class of transformations (manipulations) that can be performed on a set without altering the numerosity of the set. When called upon to explain unexpected spatial rearrangements, color changes, and item substitutions, they postulate transformations which have no effect on numerosity, such as lengthening and substitution. When probed, children typically make statements showing that they realize that these transformations do not affect numerosity." (Gelman, 1972)

In fact, by first grade, most children have a robust understanding of the difference between the types of physical transformations that do and do not change the number of objects in a set. The crucial question is How do children learn this? How do children come to classify one class of physical actions in the world as quantity-preserving transformations and another class as quantity-changing transformations? Clearly, there is no direct tutelage, unless the children happen to be unfortunate enough to have developmental psychologists as parents![15]

---

[15] And here I must thank my children (Anna, Joshua, Sophia, and Benjamin) for being willing and long-suffering sources of ideas and insights during their childhood as I subjected them to this and other forms of probes, tests, challenges, and observations … even to the point of publishing some of their behavior (Klahr, 1978). And to the extent that they thought that either Dad or they were a little odd, my apologies.

How Do Children Learn to Distinguish Between Quantity Preserving and Quantity Changing Transformations?

In order to classify the vast range of physical transformations that can be applied to small sets of discrete objects, children need three types of cognitive capacities. First, they need to be able to detect simple temporal patterns and make predictive extrapolations from them. Second, they need to be able to reliably quantify small collections of discrete objects. Third, they need to be able to parse the continuous flow of observed physical transformations in the environment into discrete temporal units having a beginning and an end. In the following paragraphs, I will summarize evidence supporting the view that children have the first two of these capacities, and I will suggest some new research paradigms that could be used to discover how and when they acquire the third capacity. The autobiographical tone of the first part of this chapter will continue, albeit in the background, in the following discussion.

Children's ability to detect and extrapolate sequential regularities

My first developmental psychology publication (Klahr & Wallace, 1970), begins as follows:

> "The ability to detect environmental regularities is a cognitive skill essential for survival. Man has a propensity to seek and a capacity to find serial patterns in such diverse areas as music, economics, and the weather. Even when no true pattern exists, humans attempt to construct one that will enable them to predict the sequence of future events." (1970, 243.)

The question we addressed was whether or not 5-year-old children could solve series completion problems. Adults' well-established ability to identify and extrapolate letter series completion problems had been modeled in a computer program (Simon & Kotovsky, 1963)[16] and our goal was to construct a computational model that could account for children's ability to solve the same general class of problems. Because we did not want to use problems that required children to have mastered the alphabet, we used a set of problems that varied in the color and orientation of simple objects, and that sometimes demanded decomposition and then reconstruction of problem attributes in order to extrapolate patterns such as those shown in Figure 5.

[16] This paper should be inducted into in the "Unheralded Landmark Papers Hall of Fame". It is the first published paper in which a theory of human performance was evaluated by directly comparing the time it took a computer model and humans to solve a set of problems that varied in difficulty. To the best of my knowledge, it has never been fully appreciated as such.

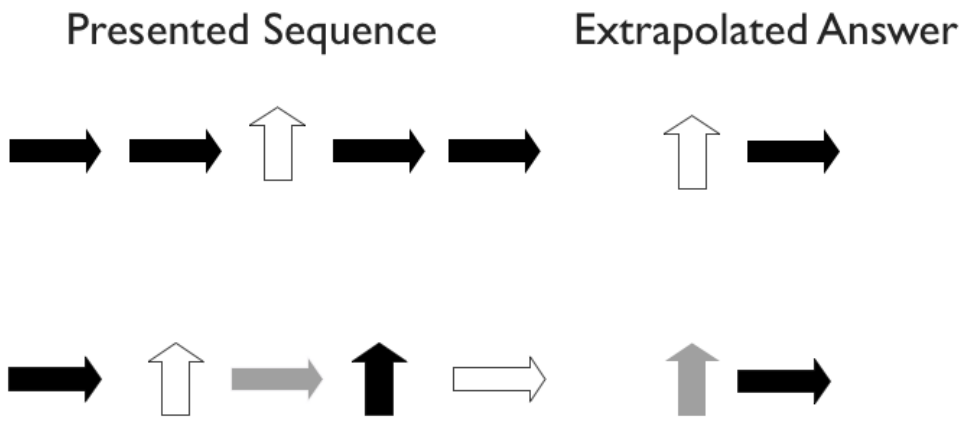

Figure 5. Series completion problems of the type used by Klahr and Wallace, 1970. In the first problem, the orientation pattern is: right, right, up, and the correct extrapolation is up, right. The color pattern is synchronized with the orientation pattern: black, black, white. In the second problem, orientation has a cycle of only 2: right-up, but color has a cycle of 3: black, white, grey. Thus the extrapolation requires the creation of a novel object (grey up).

Our investigation revealed that children could not only detect and extrapolate simple holistic patterns (such as the one shown in Figure 5a) but also decompose the dimensions of the objects, and detect and extrapolate the pattern for those dimensions[17]. That is, for problems similar to the one shown in Figure 5b, they could isolate the color pattern from the orientation pattern, and then recombine them in predicting the extrapolated item. Forty years later, the investigation of developing pattern induction capacity in children, even in infants, has become a very active research area, particularly in studies of early language acquisition (cf. Marcus, Vijayan, Rao, & Vishton, 1999; Saffran & Thiessen, 2003). There is no question that even very young children can detect, encode, and extrapolate temporal sequences of visual and auditory input.

Quantification of small sets by subitizing

Can children reliably and accurately quantify small collections of discrete objects? The literature on children's quatification abilities has identified three types of processes that are involved in encoding sets of objects and producing some kind of internal knowledge element corresponding to the size of the set: subitizing, counting, and estimation. An important developmental question is whether or not children can consistently encode sets of 1, 2 or 3 objects, before they have learned much about counting or estimation.

Evidence for subitizing as an early acquired and distict process was reported in Chi & Klahr (1975). Adults and children were presented with random dot patterns and asked to say, as fast as they could, how many were in the pattern. The results, shown in Figure 6a, show a sharp discontinuity between the reaction times for 3 and 4 items, for both children and adults. One possible alternative interpretation of these results is that they are based on a set of learned patterns, in which more objects simply allow a large number of possible cannonical patterns.

[17] These are but one of a wide range of different types of inductive problems, recently classified by Kemp & Jern (2009).

Thus 1 dot is unique, 2 dots always form a line, 3 dots always form either a triangle or a line, but 4 and more dots suddenly allow a much larger set of such forms. However, a little known, but quite important, refutation of this interpretation was provided in a study by Akin and Chase (1978). They presented adults with complex block patterns and asked them to quantify the number of blocks as fast as posible. The results, shown in Figure 6b, reveal the same abrupt change in reaction time between 3 objects and 4 objects, but these results cannot be explained by the "cannonical pattern" interpretation described above.

The question about whether subitizing –quantification of small arrays –is distinct from processes such as counting and estimation that produce internal representations of set size was a hot topic in the 70's and, in fact, remains contentious (Hannula, Räsänen, & Lehtinen, 2007; Gallistel, 2007; Le Corre & Carey, 2007, Piazza, Mechelli, Butterworth, & Price, 2002). My initial view (Klahr & Wallace, 1976) was that subitizing is indeed a distinct, and early acquired – perhaps innate, quantification process. In the early years of the debate, evidence supporting one view or another was based on the behavioral measures available at the time, but in recent years, some very sophisticated brain imaging techniques have been used to address the question, with some impressive results supporting the "subitizing is special" position.

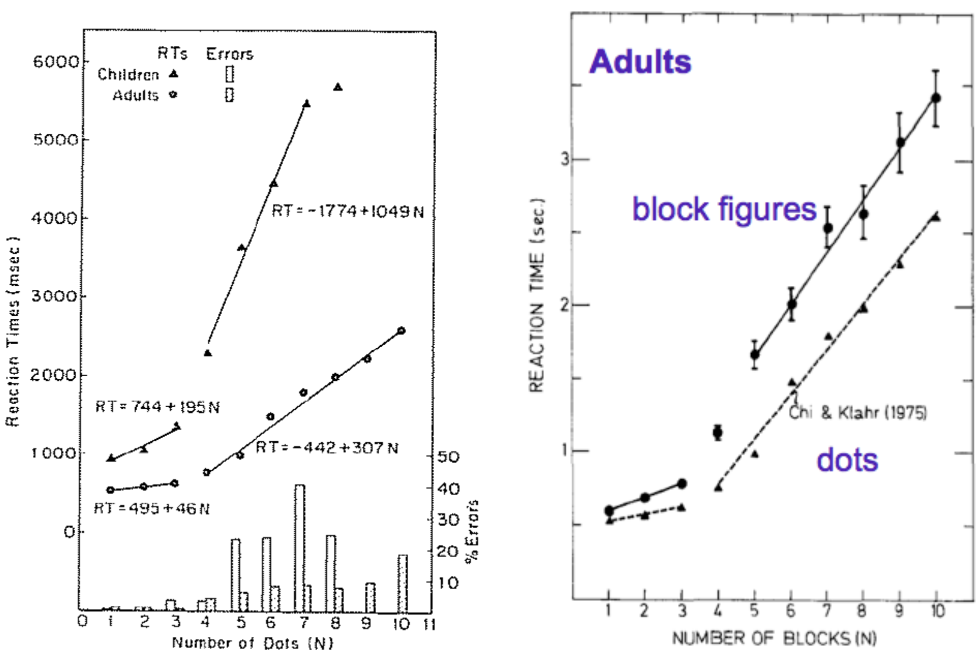

Figure 6. Subitizing and Counting RT curves. From Chi & Klahr (1975) (left), and from Aiken & Chase (1978) (right).

For example, Piazza, Giacomini, Le Bihan, and Dehaene (2003) used fMRI techniques to measure the activity of attention-related regions of the brain in a task in which adult subjects were asked to say, as rapidly as possible, the number of items in a series of displays consisting of from 1 to 7 randomly distributed items. They found a distinctively different pattern of activation between quantifying 1 to 3 items and quantifying 4 to 7 items. Their interpretation of these Blood-Oxygen-Level-Dependent (BOLD) responses and reaction time data is that different regions of the brain are activated for the higher numbers (that is, when counting, as opposed to subitizing, is occurring). Based on this evidence, and similar converging results, Dehaene (2009) concludes:

> "Although we currently have very little idea of how this system is organized at the neural level, it seems clear that a very quick and automatic grasp of the numerosities 1, 2, and 3 is part of the human intuition of numbers." (p. 244)

Parsing and segmentation of physical actions

At this point, I have argued that (1) children can detect, encode, and predict sequential regularities, and (2) that subitizing is a distinct process used to rapidly quantify small collections of discrete objects without counting. The next question that needs to be addressed is how children parse continuous actions in the world so that they can categorize actions that transform collections in ways that do or do not change their numerical amount. The answer to this question has two parts. The first part will demonstrate that–for small values of N –young children can reliably distinguish quantity-changing transformations from quantity-preserving changes. The second part will have to explain *how* such transformations are learned. That is, how a specific combination of physical actions, usually executed by a human hand, gets to be encoded as one type or another of quantity relevant transformations.

Evidence on the first part of the question is currently controversial. Some reports claim that 5-month old infants can discriminate addition from subtraction transformations (Wynn, 1998), while others claim that such competence develops slowly over the first 4 or 5 years of life (Clearfield and Mix, 1999), Huttenlocher, Jordan, and Levine, 1994; Starkey, 1992; Vilette, 1996), and that Wynn's results are based on methodological artifacts (Cohen, 2002; Cohen & Marks, 2002). There is no question that by the time children are 4 or 5, they expect what adults would call an "addition transformation" to <u>increase</u> the number of objects in a set, a "subtraction" one to <u>decrease</u> the number, and a "simple rearrangement" to leave number unchanged. However, the specific age of acquisition is not important, unless one wants to argue that this discrimination ability is innate. What *is* important is how transformational classes are learned.

What Needs to be Known to "pass" a Conservation Test?

I will address the question in two main parts. First, I will describe the kinds of knowledge components required to "DO" conservation, and the kinds of precursor knowledge components that contribute to full fledged "conservation acquisition". Although these conservation tasks appear quite simple, they actually have layer upon layer of complexity that requires the careful articulation of different aspects of the task, so that some unambiguous notation is necessary in the following exposition. After describing my notation, I will suggest an account of how those knowledge components are acquired by the child.

In Table 2, I introduce some notation for the "classic" version of the conservation of number task in which the child observes two distinct collections of discrete objects and is asked to quantify them (e.g., "how many here", "how many there?").

<u>Step 1</u> shows the following: Some kind of quantification process, Q_i operates on the first external set (set X). The subscript on Q_i indicates that the analysis is intended to cover any of the three types of quantification operators, counting – Q_c, subitizing – Q_s, or estimation – Q_e. That operator encodes collection X, and produces some internal quantitative representation of how much X is there, when quantified by operator Q_i. That *internal* representation of that amount is x_i. <u>Step 1a</u> is more specific. It shows the notation for a case in which the child subitized

collection X, thereby producing an internal knowledge element that represents the size of set X, as determined by the subitizing process. A similar process would occur for collection Y.

Step 2 is a production (Klahr, Langley, & Neches, 1987; Newell, 1990) in which the condition side is a test for whether the two internal quantitative symbols for the two collections are equivalent. If they are, then the inference is made that the *external* collections, from which the two *internal* symbols were derived, are quantitatively equivalent.

At this point, the first phase of the conservation procedure has been encoded, and the initial quantitative equivalence of two collections has been established. Next we come to the modification–or "transformation" of one of the collections. Here I introduce three kinds of generic transformations: those that increase quantity, those that decrease it, and those that preserve it. The notation is simply T_+, T_-, and T_p. The origins of the system's knowledge about how to encode and represent the observed or enacted physical actions into one of these three classes will be described below; it is, in fact, the key to the whole account of conservation acquisition being described here.

In Step 3, since we are modeling conservation rather than non-conservation in this example, we use a quantity preserving transformation on one of the collections that yields a transformed version of that collection. The notation for Step 3 indicates that a quantity-preserving transformation (T_p) was applied to collection Y, yielding collection Y'. (Think: "The set of three objects was spread out".)

Finally, in Step 4, we get to "the" conservation rule. This one says, in effect, "if you know that two collections were initially equal, and that one of them underwent an quantity preserving transformation, then you know that the quantity of the transformed collection is still equal to the untransformed one.

This final knowledge element is essential for fully understanding conservation, because it means that the system does not have to re-quantify the transformed collection in order to make a judgment about the relative magnitude of collections X and Y. To be concrete: If you put 3 cookies in a box, shut the box, and move it from you right hand to your left, you would know for sure that there are three cookies in that box. Why? Because hand shifting is a quantity-preserving transformation. It's a good thing to know!

While this decomposition may seem belabored, it suggests that young children have quite a bit to learn before they can pass conservation tests, and it may explain the otherwise surprising cases in which they have not yet acquired all of the necessary knowledge to do so. Piaget was the first to demonstrate that even when children count the same number in two collections, it doesn't mean they see them as quantitatively equivalent.

> " Aud .. counts eight pennies, says that he will be able to buy eight flowers, makes the exchange, and then cannot see that the sets are equivalent: 'There are more, because they're spread out.' These cases clearly show that perception of spatial properties carries more weight than even verbal numeration" (Piaget, 1941).

The extent to which children have this kind of knowledge, how they got it, in what contexts, for which materials and what range of numbers, and whether it could be trained, accelerated, and so on, occupied much of the field of cognitive development for about 20 years toward the end of the last century. Indeed, several of the authors in this volume, myself included, established their careers by exploring these issues.

How are the Elements of a "Conservation Rule" Acquired?

Let us assume that the cognitive system has the capacity to parse, encode, and store for further processing, the temporal sequences of external quantities and actions upon them. For example, in a data structure that Wallace and I called a "specific consistent sequence", a set of objects within the subitizing range is encoded, a specific physical transformation is observed and encoded, the collection is re-quantified, and the resulting quantitative symbols –generated by that particular quantification operator –are compared and tagged as being either the same or different. The assumption is that at the outset, transformations are encoded as quantity preserving or quantity changing, and only later are the latter types of transformations further discriminated as either addition or subtraction transformations. We called these types of knowledge elements "specific" because at this point, all that the system knows, in effect, is that in one case, if you picked up 3 coins you still had 3 coins, and in another case, if you push two dolls together you still have 2 dolls. There is no generalization here across number, transformation, or objects.

Over time, and with many such sequences, the system starts to generalize. For example, Table 3a shows a situation in which the system has encoded the fact that for a specific quantity (2 items) and a specific physical transformation (spreading), the initial and final internal quantitative representations are the same. So it learns, in effect, that "spreading doesn't change twoness".

Eventually, as shown in Table 3b, the system might discover that spreading doesn't change any of the types of quantitative symbols that subitzing is capable of producing. In other words, it will have discovered that spreading is a quantity preserving transformation, at least with respect to small collections of discrete objects. Ultimately, the system will discover, through this process of abstraction and generalization, that there are a class of quantity-preserving transformations, and another class of transformations that either increase or decrease quantity.

At this point, the system has sufficient knowledge about the relation between subitized quantities and transformations *that it doesn't have to do any requantification in order to make an inference about relative quantity, given initial quantity and the transformational type*. This knowledge is compactly represented in the following production:

$$({}_{o}x_{s})\ [T_{p}(X) \rightarrow X'\,] \quad \Rightarrow \quad {}_{o}x_{s} = {}_{n}x_{s}$$

This production rule says that:

> "*if* you know the initial quantify of a subitizable collection X (the 'old' collection), *and* you observed and encoded a quantity preserving transformation (T_p) on that collection, producing a "new" collection (X'), *then* you know, without further encoding, that the appropriate quantitative symbol for representing that quantity is the same as the previous one (${}_{o}X_{s} = {}_{n}X_{s}$).

The system gets to this state by eliminating the redundancy inherent in all the processing so far, and it creates a rule that the initial state and the transformational class are sufficient to form an expectation, a prediction if you will, about what the resultant quantity will be. At this point, the system has "acquired" conservation of quantity.

On the Basis of What Evidence is Information About Transformational Classes Noticed, Encoded, and Abstracted?

The account presented here puts the burden of acquiring conservation on children's ability to encode and classify physical transformations from the ongoing perception, encoding, storage, and classification of a huge, complex, and semi-continuous stream of visual input. How can we, as developmental scientists, gather information about how this happens? How can we document the way that children discover transformational classes in the world of continuous action sequences? I believe the answer lies in a converging portfolio of novel and powerful methodological and theoretical developments in our field: (a) statistical learning models, (b) research paradigms for discovering event segmentation capacities in adults and children, and (c) *in vivo* recording of children's observations of quantity-relevant physical transformations in the natural environment. In the following sections, I will briefly describe each of them.

Statistical learning. In the past dozen years or so, developmental psychologists have used statistical learning theory (Thiessen, 2009) to account for the way that infants encode auditory input generated by adult speech. The core idea is that there are reliable featural regularities in the continuous stream of sounds that allow the cognitive system to distinguish transitional probabilities for sequences *within* words from the transitional probabilities *between* words, and that infants can detect and use those statistical relationships between neighboring speech sounds to segment words (cf. Thiessen, & Saffran, 2007).

Do similar statistical processing mechanisms operate on visual input? More specifically, can children extract the same kind of statistical regularities from the continuous stream of physical actions they observe in the world, into classes of quantity-changing and quantity-preserving transformation? Can they derive quantitative regularities from those transformational classes? A dozen years ago, no one knew for sure. For example, at the conclusion of one of their pioneering papers on this topic, Saffran, Aslin, & Newport (1996) say:

> It remains unclear whether the statistical learning we observed is indicative of a mechanism specific to language acquisition or of a general learning mechanism applicable to a broad range of distributional analyses of environmental input. (p. 1928)

Clearly, if statistical learning mechanisms *are* sufficiently general to be independent of specific sensory modes and time scale of speech perception and segmentation, then they might be able to account for detection and segmentation of the encoding of dynamic action sequences observed by the child.

Event segmentation. The key to statistical learning is to appropriately process transitional probabilities between sequential events, and that requires an additional capacity: the segmentation of a continuous input steam into a series of discrete events. Evidence for the existence of that capacity is well-established, not only in adults (Zacks, 2004; Zacks, Kumar, Abrams, & Mehta, 2009), but also in infants (Baldwin, Baird, Saylor, & Clark, 2001). Moreover, about ten years after Saffran and colleagues' work on speech segmentation, they directly addressed the issue of how human action sequences are segmented into discrete events (Baldwin, Anderson, Saffran, & Meyer, 2008).Their study was motivated by the fact that

> … existing findings indicate that skill at detecting action segments plays a key role in processing of dynamic human activity... [but] ... the available findings have provided little insight into the specifics of how observers of dynamic action identify relevant action segments within a continuous behavior stream. That is, the mechanisms enabling adults to

extract segments from a continuous flow of activity have not been known (p. 1384).

In order to address this question, Baldwin, et al presented adults with sequences of novel and arbitrary action sequences that included mixes of high and low transition probabilities within and between segments (analogous to the research using artificial speech sounds with infants) to see whether their participants could distinguish between them (as in infant speech research). Their results show that "adults can discover sequential probabilities within dynamic intentional activity that support extraction of higher-level action segments" (p. 1401).

My earlier account of how transformational classes might be learned rests on the assumption that children can parse continuous physical actions relating to quantitative transformations into coherent units, and they can associate those units with pre- and post-transformational quantification encodings. The research on event segmentation provides clear evidence for the tenability of these assumptions. I have also argued that statistical learning theory suggests a plausible account of how the stream of visual input from quantity-relevant transformations might be segmented. An important aspect of statistical learning theory that is relevant for the kind of event segmentation I am proposing is that "many of the relations infants and adults learn involve regularities between elements that are not immediately adjacent" (Thiessen, 2009, p.37). Thus, while the purported regularities between transformations and quantification are likely to be distributed over other events, statistical learning processes could, in principle, detect and encode them. The remaining question is how developmental researchers can obtain the necessary data to further explore these claims.

In vivo recoding of transformations. What in the world do babies and young children see with respect to quantitatively-relevant action sequences? How can we discover what they see? Can we collect a corpus of everyday action encodings and then determine whether statistical learning theory can provide a plausible account of how those encodings are processed to extract quantitative transformations? Is there sufficient signal in the noise to accomplish the classification of transformations with respect to their effect on small quantities (within the subitizing range)?

In order to answer these questions, we need to extend and apply data collection paradigms that are just beginning to be developed. Put simply, if we want to know what kinds of information infants and young children encounter in the environment and what they do with it, then we need to see what they see as they encounter the real world. That is, we need to do *in vivo* research on children's encounters with quantitative aspects of the physical world. *In vivo* recording of human behavior has already proven informative in a wide range of complex human activities, ranging from observation of real scientists making real discoveries (Dunbar, 1999, 2002), to infant motor behavior (Adolph, Garciaguirre, Badaly, & Stotsky, under review), to intelligent tutors generating enormous databases as hundreds of thousands of students make second to second choices while using educational software (Baker, Corbett, & Koedinger, 2004; Romero & Ventura, 2007).

Recent work by Cicchino, Aslin, & Rakison (under review) shows how recording of infants' *in vivo* behavior can inform the theory described earlier. Using a head mounted camera, they generated a continuous record of what was in the baby's field of view (but not necessarily what the eyes were attending to), and categorized the babies' *in vivo* experiences. They found that when adults are in the baby's field of view, they are typically acting as causal agents, and babies are unlikely to observe much self-propelled action. Although Cicinno *et al* did not focus on the kind of quantitative transformations that are important for the theory I have been proposing, it is seems that most of the important transformational classifications will come from situations in which adults are indeed the agents of change with respect to small quantities.

The opportunity for *in vivo* research on what children see in the natural world (as opposed to the psychologist's laboratory) has been substantially facilitated by the recent development of a sophisticated but very light weight eye tracker that can be mounted on infants' heads as they negotiate their everyday environment – rather than sit strapped into seats in the researcher's lab -- to reveal exactly where in the scene the baby is looking (Franchak, Kretch, Soska, Babcock, & Adolph, 2010). In the first publication on this new technology, Franchak, *et al* report a very high proportion of infants' fixations being directed to the manipulation of objects in mothers' hands. While this study was not designed to focus on quantification or on quantity-related transformations, it is clear that such studies could be mounted. It is this level of dense data recording that will enable us to collect the necessary information about the frequency, reliability, and grain size of children's exposure to the kinds of data necessary to assess the hypothesized processes sketched in Table 2.

Concluding comments

This chapter is an odd bird, being a mix of two things that scientists are trained to avoid: excessive personalization and highly speculative theorizing. Fortunately for me, the organizers and editors of this Festschrift have encouraged and indulged me in this regard, and I hope that the reader finds the final product interesting. I mean that with respect to both parts of the chapter. I hope that the first part motivates readers to think about their own emotional attraction to their work, and the developmental paths to their current identification as researchers. I also hope that some readers will find my scientific speculations in the final part of the chapter sufficiently intriguing to expand and implement the suggestions made here about how to advance our understanding of quantitative development.

Acknowledgements

My deepest thanks to organizers and editors of this Festschrift: Jeff Shrager, Kevin Dunbar, and -- most of all -- Sharon Carver, for the energy, generosity, respect, and affection that they put into the project. I am also indebted to Audrey Russo, the administrative assistant for the Symposium, for her creativity, perseverance , enthusiasm, and assistance not only with respect to this chapter, but in all my professional activities. Finally, but certainly not least, I offer my thanks and love to my wife for always finding the right balance of curiosity, challenge, and sustenance for both the affective and the cognitive aspects of my life

References

Adolph, K. E., Garciaguirre, J. S., Badaly, D., & Stotsky, R. (under review). How infants learn to walk: 15,000 steps and 100 falls per day. (*Psychological Science*)

Akin, O. & Chase, W. (1978) Quantification of Three-Dimensional Structures. *Journal of Experimental Psychology: Human Perception and Performance*. 4, 397-410.

Baker, R., Corbett, A., & Koedinger, K. (2004). Detecting student misuse of intelligent tutoring systems. In *Intelligent tutoring systems* (pp. 531–540)

Baldwin, D. A., Baird, J. A., Saylor, M. & Clark, M. A. (2001). Infants parse dynamic action. *Child Development*, *72*, 708–717.

Baldwin. D, Andersson, A., Saffran, J. & Meyer, M. (2008) Segmenting dynamic human action via statistical structure. *Cognition*, 106, 1382–1407.

Carver, S. M. & Klahr D. (Eds.) (2001) *Cognition and Instruction: 25 years of progress*. Mahwah, NJ: Lawrence Erlbaum Associates.

Clearfield, M. W. & Mix, K. S. (1999, April). Infants use contour length – not number – to discriminate small visual sets. Poster presented at the biennial meeting of the Society for Research in Child Development, Albuquerque, NM.

Cohen, L. B. & Marks, K. S. (2002). How infants process addition and subtraction events. *Developmental Science*. *5*, 186-201.

Cohen, L. B. (2002). Extraordinary claims require extraordinary controls. *Developmental Science, 5,* 211-212.

Cicchino, J. B., Aslin, R. N., & Rakison, D. H. (under review). Infants' visual experiences shape their representations of causal and self-propelled motion. Cognition

Chi, M.T., & Klahr, D. (1975). Span and rate of apprehension in children and adults. *Journal of Experimental Child Psychology*, 19(3), 434-439.

Dehaene, S. (2009) ,Origins of Mathematical Intuitions: The Case of Arithmetic. *The Year in Cognitive Neuroscience 2009*. Annals of the New York Academy of Science, 1156: 232–259.

Dunbar, K. (1999). The Scientist *In Vivo*: How scientists think and reason in the laboratory. In Magnani, L., Nersessian, N., & Thagard, P. *Model-based reasoning in scientific discovery*. Plenum Press.

Dunbar, K. (2002). Science as Category: Implications of InVivo Science for theories of cognitive development, scientific discovery, and the nature of science. In S. Stich & P. Carruthers (Eds.) *Cognitive Models of Science*. Cambridge University Press.

Franchak, J. M., Kretch, K. S., Soska, K. C., Babcock, J. S., & Adolph, K. E. (2010). Head-mounted eye-tracking in infants' natural interactions: A new method. *Proceedings of the 2010 Symposium on Eye Tracking Research and Applications*, Austin, TX.

Gallistel, C. R. (2007). Commentary on Le Corre & Carey, *Cognition*, 105, 439–445.

Gelman, R. (1972) Logical capacity of very young children: number invariance rules, *Child Development,* **43** 75–90.

Hannula*, M. M.*, Räsänen, P. & Lehtinen E. (2007). Development of counting skills: Role of spontaneous focusing on numerosity and subitizing-based enumeration. *Mathematical Thinking and Learning,* 9 (1), 51-57.

Huttenlocher, J., Jordan, N. C., & Levine, S. C. (1994). A mental model for early arithmetic. *Journal of Experimental Psychology, 123*, 284–296

Kemp, C. & Jern, A. (2009) A taxonomy of inductive problems <????>

Klahr, D. (1969a). Statistical significance of Kruskal's nonmetric multidimensional scaling technique. *Psychometrika, 34*, 190-204.

Klahr, D. (1969b). Decision making in a complex environment. *Management Science, 15*, 595-618.

Klahr, D. (1973). An information processing approach to the study of cognitive development. In A. Pick (Ed.), *Minnesota symposia on child psychology, Vol. 7* . Minneapolis: University of Minnesota Press

Klahr, D. (Ed.). (1976). *Cognition and instruction*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Klahr, D. (1978). Goal formation, planning, and learning by pre-school problem solvers, or: `My socks are in the dryer'. In R.S. Siegler (Ed.), *Children's thinking: What develops?* (pp. 181-212). Hillsdale, NJ: Lawrence Erlbaum Associates.

Klahr, D. (1984). Transition processes in quantitative development. In R. Sternberg (Ed.), *Mechanisms of Cognitive Development* . (pp. 102-139). San Francisco: W. H. Freeman & Co.

Klahr, D. (2004). Encounters with the force of Herbert A. Simon. In Augier, M. & March, J. G. (Eds.) *Models of a Man: Essays in Memory of Herbert A. Simon.* Cambridge, MA : MIT Press.

Klahr, D., Langley, P., & Neches, R. (Eds.). (1987). *Production system models of learning and development.* Cambridge, MA: MIT Press.

Klahr, D., & Wallace, J. G. (1970). The development of serial completion strategies: An information processing analysis. *British Journal of Psychology, 61*, 243-257.

Klahr, D., & Wallace, J. G. (1970). An information processing analysis of some Piagetian experimental tasks. *Cognitive Psychology, 1*, 358-387.

Klahr, D., Fay, A.L., & Dunbar, K. (1993) Developmental differences in experimental heuristics. *Cognitive Psychology, 25.*, 111-146.

Krumboltz, J. D. (Ed.) (1966) *Learning and the Educational Process* . Chicago: Rand McNally & Co.

Kruskal, J. B. (1963) Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29, 1964

Le Corre, M. & Carey, S. (2007) One, two, three, four, nothing more: An investigation of the conceptual sources of the verbal counting principles. *Cognition*, 105, 395–438

Marcus, G. F., Vijayan, S., Rao, S. B., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, 283, 77–80.

Newell, A. (1990) *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press, 1990

Piazza, M., Giacomini, E., Le Bihan, D., & Dehaene, S. (2003). Single-trial classification of parallel pre-attentive and serial attentive processes using functional magnetic resonance imaging. *Proc. R Soc. Lond. B Biol. Sci.*, *270*(1521), 1237– 1245.

Piazza, M, Mechelli, A. Butterworth, B. & Price, C. (2002). Are subitizing and counting implemented as separate or functionally overlapping processes? *NeuroImage*, 15:(2), 435-446.

Reitman, W. R. (1964) Heuristic decision procedures, open constraints, and the structure of ill-defined problems. In M. W. Shelly and G. L. Bryan (Eds.), *Optimality and human judgment.* New York: Wiley, 1964.

Reitman, W. R. (1965) *Cognition and Thought*. New York: Wiley.

Romero, C. & Ventura, S. (2007) Educational data mining: A survey from 1995 to 2005. Expert Systems with Applications 33 (2007) 135–146

Saffran, J.R., & Thiessen, E.D. (2003). Pattern induction by infant language learners. *Developmental Psychology,* 39, 484-494.

Saffran, J. R., Aslin, R. N. & Newport, E. L. (1996) Statistical learning by 8-month-old infants. *Science*, 274, 1926 -1928.

Schunn, C. D., & Klahr, D. (1992). Complexity management in a discovery task. Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society.

Schunn, C. D., & Klahr, D. (1996). The problem of problem spaces: When and how to go beyond a 2-space model of scientific discovery. In the Proceedings of the 18th Annual Conference of the Cognitive Science Society.

Starkey, P. (1992). The early development of numerical reasoning. *Cognition, 43*, 93–126.

Thiessen, E. D. (2009). Statistical learning. In E. Bavin (Ed.), Cambridge Handbook of Child Language, pp. 35-50. Cambridge: Cambridge University Press.

Thiessen, E. D., & Saffran, J. R. (2007). Learning to learn: Acquisition of stress-based strategies for word segmentation. *Language Learning and Development*, 3, 75-102

Vilette, B. (1996). De la "proto-arithmetiques" aux connaissances additive et soustractives [From "proto-arithmetic" to additive and subtractive knowledge]. *Revue de Psychologie de l'éducation, 3*, 25–43.

Wakeley, A., Rivera, S., & Langer, J. (2000) Can Young Infants Add and Subtract? *Child Development*, 71, 1525–1534.

Wynn, K. (1992). Addition and subtraction in human infants. *Nature, 358*, 749–750.

Zacks, J. (2004). Using movement and intentions to understand simple events. *Cognitive Science*, 28, 979–1008.

Zacks, J. M., Kumar, S., & Abrams, R. A., Mehta, R. (2009). Using movement and intentions to understand human activity. *Cognition*, 112, 201-216.

Table 1. Personal shaping forces, contexts, and lessons learned

Career point	Topic	Context	Forces	Lesson Learned
Pre-teen	Category formation and parental approval: logic and love	Helping my father in his business by sorting watch parts	Parental approbation and appreciation.	Classification requires creativity. I can do it. It is valued.
Teen & College	Abstract representations of reality: surveys and maps.	After-school job as Surveyor's Assistant	Intellectual appreciation of "real world" measurement process.	Measurement processes abstract, refine, simplify, but also create knowledge.
Post-College	Knowledge driven search trumps trial and error	Programmer at NORAD, tracking satellites and missiles	Complex computations on limited computing devices	Advances arise from ingenious problem formulations not brute force computation
Early Grad School	Serendipity at Stanford	Summer conference on “Learning and the Educational Process” at Stanford	Exposure and introduction to an entirely new set of problems: Piaget's theory of cognitive development.	"Secret Weapons" can be brought to bear on well-established problems.

Table 2 : Processing Steps for Equivalence Conservation

1. $Q_i(X) \longrightarrow x_i \qquad Q_i(Y) \longrightarrow y_i$	A quantification operator encodes the external collection (X) and produces an internal symbol indicating the amount and the type of quantifier that produced that symbol. Ditto for external collection Y.
1a. $Q_s(X) \longrightarrow x_s \qquad Q_s(Y) \longrightarrow y_s$	In this case, the operator is Subitizing.
2. $(x_s = y_s) \longrightarrow (X \overset{Q}{=} Y)$	If the subitizing symbols from the two collections are the same, then you know that those external collections are quantitatively equivalent.
3. $T_p(Y) \longrightarrow (Y')$	Collection Y undergoes a "perceptual" transformation (i.e., NOT T_+ or T_-)
4. $\{X \overset{Q}{=} Y\}$ & $T_p(Y) \longrightarrow (Y') \Longrightarrow \{X \overset{Q}{=} Y'\}$	Now you know that If two collections are initially quantitatively equivalent, and one undergoes a perceptual transformation (from Y to Y') then you know that collection X is quantitatively equivalent to collection Y'.

(a) Generalization over objects (for a specific subitizing symbol)

$$\left.\begin{array}{l}\text{two dolls}\\ \text{two cookies}\\ \text{two fingers}\end{array}\right\} \ldots \text{ spread apart } \ldots \left\{\begin{array}{l}\text{two dolls}\\ \text{two cookies}\\ \text{two fingers}\end{array}\right.$$

$${}_{o}2_{s} \quad \ldots \quad \text{spreading} \quad \ldots \quad {}_{n}2_{s} \quad \ldots \quad {}_{o}2_{s} = {}_{n}2_{s}$$

(b) Generalization over quantitative symbols

$$\left.\begin{array}{l}{}_{o}2_{s}\\ {}_{o}1_{s}\\ {}_{o}3_{s}\end{array}\right\} \ldots \text{ spreading } \ldots \left\{\begin{array}{l}{}_{n}2_{s}\\ {}_{n}1_{s}\\ {}_{n}3_{s}\end{array}\right.$$

$$\text{any } {}_{o}x_{s} \quad \ldots \quad \text{spreading} \quad \ldots \quad {}_{n}x_{s} \ldots {}_{o}x_{s} = {}_{n}x_{s}$$

(c) Generalization over transformations

$${}_{o}x_{s} \quad \ldots \left\{\begin{array}{c}\text{spreading}\\ \text{rotating}\\ \text{compressing}\end{array}\right\} \ldots \quad {}_{n}x_{s} \quad \ldots \quad {}_{o}x_{s} = {}_{n}x_{s}$$

(d) Common consistent sequences:

$${}_{o}x_{s} \ \ldots T_{p}(X) \longrightarrow X' \ldots \ x_{s} \ \ldots \ x_{o} = x_{s}$$

$${}_{o}x_{s} \ \ldots T_{+/-}(X) \longrightarrow X' \ldots \ x_{s} \ \ldots \ x_{o} \neq x_{s}$$