

Discussion between Steven Hayes and Bill Powers 2011

Hello, all --

Attached is a nice article by Fred Nichols (<http://www.nickols.us/ManageYOP.pdf>) who is a management consultant and a long-time CSG (control system group) member. It has a bearing on the difference between ACT (Steve Hayes' therapy system) and MOL (the method of levels, which goes with PCT). People have to do their own controlling, is the general theme.

I have long claimed that while all therapies prove to have some degree of success, they are successful only to the degree that they incorporate the same principles and processes that constitute MOL. This now includes ACT, which more than any other system of therapy I have seen explicitly recognizes the basic principles and phenomena that underlie MOL.

My colleague David Goldstein, after hearing my description of Steve's approach, found a Wiki article:

http://en.wikipedia.org/wiki/Acceptance_and_commitment_therapy

Here is an interesting passage from it:

"ACT helps the individual get in contact with a transcendent sense of self known as "self-as-context" - the you that is always there observing and experiencing and yet distinct from one's thoughts, feelings, sensations, and memories."

I discovered this sense of self during discussions and explorations with a friend, Kirk Sattley, somewhere around 1953 or 1954 -- actually six or seven years before the first publication on "feedback theory" in 1960, which became PCT. The method of levels came out of those explorations, in time for my 1973 book, but the editors felt that that chapter (along with the one on emotion) did not have enough to do with the rest of the book and chopped it out.

Since then, this entity has come to be called the "Observer Self," the one that is "aware." All other perceptions result simply from the actions of neural networks and can be treated as objects of awareness. The world we experience, this seems to imply, is how neural signals look to the Observer Self.

For me, this is the core idea of MOL, and of ACT as it was described to me. I think it was really discovered a long time ago in the Far East, but because the idea is spooky and mysterious, science has generally ignored the phenomena, despite the fact that they are quite reliably reproducible. MOL and ACT (and "mindfulness" and other approaches) are, I think, establishing a foothold for accepting this concept as an approved subject for scientific examination.

What I am seeing within the MOL community is a lot of conflicts that arise from the previous orientations of people trying to learn the MOL approach. Conventional therapists find it difficult to give up the problem-solving, prescriptive, analytical, doctor-knows-best approach and simply help a client explore the mental territory. I can see that a behaviorist would have a problem with giving up the "prediction and control" or "analysis of behavior" attitudes. But whatever our origins, it is the recognition of the Observer Self that gives us all a common playing field and that will eventually help each of us strip away the remnants of older approaches and focus on the central processes that are the real core of successful therapy. Control theory is just a useful idea. So is analysis of behavior. Useful ideas change with time and experience, especially if one can avoid identifying with them and find a more elevated position from which to view them. We need to find a way to formulate the highest-level goals, as Steve might put it, so their achievement is not frustrated by lower-level disagreements. Conflicts are resolved by going up a level (or a few levels), not by out-arguing someone. I think the common goals that draw us onward are at a high enough level to make what we give up seem a relatively small loss.

So, where do we go from here?

Best,

Bill

I'm a baby on the PCT front but the ACT folks have tried to take this "self as context" idea down into the basement to get the elves to work on it. We are busy doing basic work on deictic framing to see if we can understand how this sense of self emerges.

I've attached a recent review of ACT and its outcomes and processes. Is there a similar empirical / conceptual review of MOL that I might be sent?

- S
-

ACT Processes and Outcomes BRAT 2006.pdf

<http://www.sciencedirect.com/science/article/pii/S0005796705002147>

Deictics-Weil-PsychRecord--figures 1-3.doc

Deictics-Weil-PsychRecord 2 18 2010 submitted.doc

[For other work by Weil see:

<http://gradworks.umi.com/32/39/3239878.html>]

.....

Hi, Steve --

Good, I thought you might reply. It is pretty hard to contemplate the possibility that somebody else's idea might contain something one has missed, but I'm trying and I think I can do it. I expect that you can, too.

OK, I've signed up and paid a little bit of dues and am now officially a neophyte in ACBS. I'm glad you kept the price down, being retired and not yet having received my Nobel Prize (can't think what the delay is).

I had a look at a brief description of "diectic framing", and while I clearly have to read more about it, the first resonances suggest that it has something to do with what we call levels of perception and control. Not the specific ones I have proposed, which remain to be seriously tested as universals (so far, N is somewhere between 1 and 5), but simply levels, whatever they prove to be for an individual. Higher-level perceptions are functions of sets of lower-level perceptions, and control systems for higher-level perceptions operate by varying the reference signals for (setting the goals for) lower-level systems. Actions, or behavior observable by others, are generated only at the lowest level, the spinal control systems formerly known as reflexes. That's the basic structure of PCT. I have lately discovered that control engineers call this a "cascaded system." I thought I had invented it. Too bad.

My mental model of what is going on here isn't far advanced and it makes no distinction between science and science fiction. But it does make a degree of experiential sense.

What seems to me to be happening is that awareness, the perceptual input function of the Observer Self, receives information from the neural hierarchy of perception at variable levels of organization and within various systems at the levels in question. The "flashlight in the cathedral" is one image I use to communicate this idea. As awareness moves from one place in the hierarchy to another (you would probably say "from one context to another"), the content of consciousness

changes, and the individual experiences the world as if from the point of view of the neural systems lit up by the "flashlight." If you are conscious of (third-order) objects, for example, the entire world, in all sensory modalities, seems to be made exclusively of objects (configurations like chairs and musical chords). If you look very closely to see what objects are made of that are not just smaller objects, you find that the world has changed into a collection of sensations: edges, pressures, colors, tones, smells, and so on (second order). The world seems to take on the characteristics of the level of perception currently associated with awareness. If you look at the world as a collection of sensations, objects, transitions, and events, you may well find yourself seeing in these collections a rather large assortment of relationships such as in, on, above, behind, because of, higher, louder, nicer.... or even more complex ones. At a higher level you see categories of relationships which you name and refer to by words (like "in", "on", "above" ...). The symbols used as names are parts of sequences, ordered lists which are parts of logical (rule-driven) processes, which are organized to maintain principles, which are the elements of system concepts.

"Values" are what we call "reference signals" or "reference conditions." They specify a perception as a blueprint specifies a house, and behavior is organized -- when everything is working right -- to make perception approach and then remain close to the reference condition being set by a still-higher-order system. The term "value" would refer to reference signals (and perceptions) fairly high in the hierarchy; at lower levels a reference signal is just a target for details like sensed forces or positions or velocities.

The main thing that can go wrong in this hierarchy of control (other than physical damage) is that higher systems sharing the use of some set of lower systems fall into conflict with each other: for one system to make its own perception match its reference condition is to prevent the system on the other side of the conflict from doing the same. The person ends up trying to do one thing and its opposite at the same time, or to do and not do the same thing. This is what you call "futile attempts to wage war against their own inner lives." There is nothing wrong with the control systems involved except that as they try to behave, they are disturbing another system which, in counteracting the disturbance, returns the favor. Neither system can control under those conditions, and those systems are lost to the hierarchy as useful subsystems. The better both systems could control by themselves, the worse is the result.

I think this may be what you refer to obliquely when you say "control is the problem." I don't think you think that the ability to make your experiences be the way you want them to be is a problem, but simply that one human controller can't control another one arbitrarily and attempts to do so create problems. At first, many people are put off by the phrase "control theory" because they think it's all about how one person can control another one, or how people can use will power to control (i.e., wage war against) themselves. But PCT shows how a hierarchy of control systems can be organized so that hundreds of control systems can come to operate without any conflict between them at all -- and it explains why the results are so bad when control systems in one organism try to incorporate another one's behavior into the set of variables they control. In the latter case, arbitrary control definitely IS the problem.

Your approach is basically organized around analysis of the elements of behavior, and while there are some loose ends to clear up there, I think we can benefit from your findings. Identifying controlled variables is a big part of PCT methodology, and we haven't done much of it, while you seem to have done a lot. Your approach to therapy seems to involve considerably more intervention and prescription than our MOL does, but that may be just a surface appearance and I can't imagine that you have failed to notice how such things interrupt the flow of therapy. But we can have some nice arguments about that if required.

Probably the biggest difference between ACT and PCT is that you are focused on the *what* of behavior while we focus more (at least I do) on the *how* -- the actual mechanisms. I like to understand how things work, and PCT represents my attempt to understand how organisms work. At the moment, I can't see any reason why we can't trade knowledge. Undoubtedly there will be differences, but my intention is that we track down their origins and get rid of them. If you have the

same intention, we shouldn't have any major problems.

Best,

Bill

Hello, Steve --

I've been looking at the last post from you for weeks now, trying to figure out what I want to say about it and the materials I followed up by reading. The difficulty is that we have both developed very similar basic understandings of how consciousness relates to the operation of the brain, and we are ideal people for developing a strong rivalry. Or that is what I thought at first.

Today I read some of your Psychology Today blogs.

(for the cc'd audience,

<http://www.psychologytoday.com/blog/get-out-your-mind>

)

What you talk about in these moving and insightful short sermons is the kind of attitude and interpretation of experience that my own approach, the method of levels (MOL), also leads to. Mine is not pronounced as it looks: just say M. O. L. In addition, you have introduced ideas of which I have been a follower since reading A. E. Van Vogt's "The world of null-A" and discovering that the pithy quotes scattered through the chapters actually came from a real book, Korzybski's "Science and Sanity," which I quickly acquired and devoured, at the age of 15 or 16. I immediately understood what "The map is not the territory, the word is not the object" meant, and some years later, off that springboard, came to understand that the word is not the perception, either, nor is the observer the thing observed. Somewhere in that progression, PCT began to take shape, and eventually those two threads came together in the form of MOL.

Of course I wondered if you, too, had stumbled across Korzybski, but regardless of that our foundational concepts -- you might even say Second Foundational concepts -- are more compatible than different.

The differences will eventually have to be laid out, discussed, and settled. But I take heart from the mantra at the end of your posts: "Love isn't everything, it's the only thing." Love assumes good will beneath everything, and if we try to stick to that we should manage to find common ground.

What we each claim as the science behind our approaches is very different, so let's put that aside for a while. I want to raise a question of technique.

As I read your blogs there were many moments of recognition and agreement, when I admired your way of putting something and empathized with the experiences that led you this way. At the same time, I was thinking "This isn't how I do MOL with people." It's more like the way I would discuss the impact of MOL with people who had already been to these places and had these realizations.

My main approach with MOL is to try to lead the client (though I don't have any actual clients) to arrive at the main conclusions in the client's own words and out of the client's own experiences. The way the other person does this is a constant source of surprise and amazement to me. If I can only get the person to pause and talk about the background thoughts that come up, the person will soon tell me what the problem is, and soon afterward, a solution. The one place where I get more explicitly directive is when a conflict shows up, which is quite frequently. Then I bias my questions

toward bringing the two (or more) sides of the conflict into attention for discussion, and go back and forth between them until the client says something to indicate that both sides are in attention at the same time and that we have found the core of the conflict. At that point I have done what I can, and the usual sequel is that the conflict simply dissolves, and we can get back to going up levels.

I have to explain that my "clients" are almost exclusively people learning to do the method of levels, who have volunteered to help me put on a demonstration of the method, often in front of a friendly audience of colleagues. So my mean time with a client is about 15 minutes at most. In the beginning, while working out the methods, I spent many hours with my friend Kirk Sattley and a few others who fled from dianetics when I did (!), switching roles back and forth, so I am not without extended experience in the chair or on the couch, but a therapist I am not. On the other hand, most of the demonstrations turn into real sessions in which real problems come up and -- as shown by follow-up enquiries -- have lasting resolutions. As you commented about ACT, the results come very quickly when you do all the necessary things and avoid the unnecessary or counterproductive things like giving advice, trying to solve the client's problem instead of letting the client do it, or diagnosing the client (or even thinking of diagnosing the client).

So there will you detect a difference in our approaches? Perhaps.

Here's another. You advise people to accept themselves, whether good or bad, and simply observe. I get clients to do that by asking questions. If the client says "I really have no social skills at all," I begin by asking something like "Does that bother you?" And then we talk for a while about why it is bothersome, usually coming across several conflicts as we go. I don't tell them they shouldn't be bothered, or explain the benefits of acceptance. I just get them to resolve the conflicts about social skills, so their self-critical attitude goes away by itself. They get to the same place your approach leads to, I think, but without their being told to go there.

I don't actually know which way is quicker or more effective, and I'm sure you have a lot more successes than non-successes, so there wouldn't be any dramatic differences. But I am trying to develop a minimalist therapy in which only what helps is done, and nothing is done that works in the other direction. Most therapies I have learned about do a great deal that isn't necessary, like talking to chairs or rolling the eyeballs. They perhaps do even more that interrupts progress.

So how are the vibes at this point?

Best,

Bill

Hi, Steve --

I've started reading the "Talking ACT" book, and what you actually do in a session is similar enough to MOL (other than the therapist being a bit more in control of the session in ACT) that there is hardly anything to have an interesting argument about.

Contrary to what I said before, that really leaves us mostly with the background science to talk about. What we imagine to be going on during a session is different, because we have different, though not totally different, conceptions of how behavior works. As I read the first few pages, through the first session transcripts, I kept thinking that I might talk about the same phenomena but I would explain them differently, which would lead to asking different questions of the client and perhaps steering the conversation in a somewhat different direction. But after half a dozen instances of this, I realized that the specific comments I wanted to make all had the same theme: the theory behind the descriptions and actions. At least those are the words that occur to me now; the idea will probably get clearer as we move along.

What you and I think of in connection with the word "model" are different things. Your six core processes involve phenomena that I would say need models to explain them; they are not models in themselves. A model enables you to say what process generates a phenomenon such as acceptance, defusion, and so on. In the case of acceptance-avoidance, what makes the difference is a reference level for some perception that is set either high or low, so that presence of the variable being experienced is wanted or rejected. What makes the difference between acceptance and rejection is the higher system that (in PCT) controls some higher-level perception by means of varying the reference level for acceptance. The question then becomes, acceptance or rejection for what purpose? In MOL we would help the client look for the higher-level control process which answers that question. People don't accept or reject experiences just to be accepting or rejecting them; there's always a higher goal to explain what they do. The therapist doesn't know what that goal is for any one person, but the therapist knows how to ask questions that will help the client explore until an answer is discovered.

What I'm saying is not that your model is no good, but that PCT offers another layer to slip beneath it that will provide a more secure foundation. Here's an example from the Dialog:

=====
An example of an acceptance exercise could involve imagining a distressing experience, such as an argument with a friend, and working to sit with whatever feelings occur. An example of an acceptance metaphor is the two-games metaphor.

[DIALOG]

T: It is sort of like you are playing against a professional team; there is just no way you can win.

C: Yes, that is what it feels like.

T: What would you get if you won that game?

=====
The therapist's question is the same kind of question that an MOL therapist would ask (if the client doesn't spontaneously provide the information as usually happens). In terms of the PCT model, the therapist is asking "What higher-level goal would become possible to achieve if you achieved the lower-level goal?" In other words, the therapist is inviting a statement about higher levels of control, and he gets it.

=====
C: I feel like once I get this part of my life under control, I can do the things I have been missing like dating, changing my job, and being more active.
=====

AT this point MOL diverges from ACT. The ACT therapist says

=====
T: Here is my offer. There is another game over here that most people do not pay attention to [puts out other hand]. It is similar to the first game in some ways, but also different in other very important ways. To begin with, this game is fair. The more you put into it, the more you generally get out of it. Most importantly, instead of playing for control of your thoughts, you play for those things that you really want in life. Instead of getting your emotions under control before you move forward, what if we just started moving forward?
=====

The ACT therapist, it seems to me, has dropped into problem-solving mode and is going to make some suggestions about how the client might change the game. This may in fact (with luck) prove useful, but in MOL we would follow a different course that I think is more likely to result in immediate improvements.

What we are hearing from C is that he is struggling with a thought or feeling, and if he could resolve the struggle, he knows what ends that he desires would be easier to reach -- dating, job, activity. So the PCT therapist immediately asks himself, "Why is he still having this problem? Why doesn't he just reorganize a little and resolve it?" That raises a red flag, so the next question then considered is "What is the conflict?"

Against what is C struggling? In its simplest form, one possibility is that he is mad at the friend he has had an argument with and still wants to win it, but he doesn't want to lose the friend, so he doesn't want to win it. That's the core form of a conflict, a logical contradiction: I want to have condition X (for one set of reasons), and at the same time I want to have the condition not-X (for a different, and seemingly just as good, set of reasons). The result is disorganization and inaction or vacillation.

A basic PCT idea is that about the worst non-physiological thing that can happen to a brain is for two properly functioning and competent control systems to start telling lower systems to accomplish two contradictory goals at the same time. The more skillful the control systems are, the more extreme is the resulting escalation of the conflict. You suggest elsewhere that it's a war, and I would agree it can turn into that, a war against yourself.

So C is struggling with himself because there is some fairly important perception that he wants to be in two different states at the same time, which is impossible. He is keeping himself in conflict by focusing on different ways to have a war, when he should be focusing on what is causing the conflict: the level of organization that is setting the contradictory goals.

Another PCT precept is that "reorganization follows awareness." This is why we want the client to become aware as if (or really) from a higher level of perception. The attention is stuck at the level of the conflict, so the client is trying to reorganize at the level of how to do two contradictory things at once. The reorganizing process won't resolve the conflict that way. Awareness has to be directed toward the higher levels of perception and action that are setting the contradictory goals so those control systems will be reorganized and remove the reason for the conflict.

I think that about all that needs to be done is to get the attention up a level, where the causes of the conflict reside, and keep it there for a while. The rest will happen because of the innate capacity for reorganization that the client has, only now applied where it can do some good.

This is enough for one dose, but there's one more pertinent factor. How can you tell when therapy is finished?

This is my personal answer, imperfectly validated for anyone else but definitely OK for me. Therapy is finished when you can go through an MOL session tracing out the structure of levels of control until you reach a level where there is nothing further to discuss -- without encountering a single conflict that doesn't dissolve immediately upon being noticed. Conflicts keep coming up, but once you learn to recognize them and step back (or up) to a more helpful point of view, they are very quickly resolved without needing any external help.

I haven't specifically mentioned any differences in the science bases supporting our approaches, but I think you have probably noticed the lack of some terminology and the presence of others. We can get to that later if you wish.

Best,

Bill

ACT folks would agree that all of these processes (acceptance, defusion, etc) require an explanation.
We are trying to build that from the bottom up by research on basic behavioral principles and cognitive processes.

For example, I've attached on of the early papers on experiential avoidance (184)

Then I've attached one of the early papers on Relational Frame Theory (144)

Finally an example of experimental studies linking the two (Nic Hooper study)

We have work like that in all six areas of the model

I've also attached a paper that describes the strategy itself (ACT and CBS)

When is therapy done:

I'd say it was done when the person had learned how to have life itself serve as the teacher.

The way we think about it, experience has been blocked from its natural role as a guide to growth and development due to these inflexibility processes. We help people get out of their way so that experience begins to shape successful action linked to chosen values. Once that happens with some regularity, therapy is over and life is the therapists.

People can come back for boosters but people are not broken and it is not our job to fix them or to live their lives for them. Another way to say it is that our job is skillful liberation, not "cure."

- S

184 Experiential avoidance JCCP 1996.pdf

<http://psycnet.apa.org/journals/ccp/64/6/1152/>

144 Verbal Relations AP 1992.pdf

<http://psycnet.apa.org/psycinfo/1993-12957-001>

Nic Hooper study on thought suppression.pdf

<http://www.springerlink.com/content/v18422p423146415/>

Hi, Steve --

I've enlarged the audience a bit to include neuroscientist Henry Yin, and my collaborator on the programs in my latest book, Bruce Abbott, who I think would not object to being described as a behaviorist. Or maybe a recovering behaviorist.

The "Behavior therapy" paper (attached) gives me an even better focus on ACT. Your objections to CBT are the same as mine. Beck is basically just trying to get clients to think like logical, rational, scientists. There, see? There's nothing to be afraid of so stop being anxious. Not that any competent therapist would actually put it that way, CBT or no CBT.

Furthermore, your descriptions of and comments on the basic phenomena of awareness are the same ones I have noticed, and your exercises are like mine down to asking clients to examine

thoughts from a non-thought point of view. I strongly suspect that we are far from the first to notice such phenomena or to devise such exercises.

So at the operational level of discourse there is little to disagree about in those regards. In fact it is encouraging to see that two people can agree so well about the phenomena even though they arrived at their descriptions independently and have some serious differences when it comes to explaining them. We can at least agree that we are competent observers and claim that we are observing something real.

You speak about science being in a continuous state of change, and I think we are looking at another instance of that principle. It comes down to PCT versus radical behaviorism and your later developments from that base. We each have a tremendous investment in our own interpretations of the evidence, a lot of people backing us with their support and belief, and a lot of claims in print. That is a bad position for a scientist to find himself occupying. All the hoopla and excitement is fun, but such enthusiasm makes it very difficult to change any ideas. People are counting on us to give them True Sayings, and if we were to remark that maybe some of those Sayings are no longer true, our own followers would object loudly. The solid ground would give way beneath their feet. This has already happened to me when I have changed my mind about some part of PCT, like moving the sequence-control level from level 5 into a new level 8. I don't doubt that it has happened to you, too. I have spent years trying to persuade people not to memorize the levels I proposed or take them too seriously, and to propose their own. I was really just trying to work out how one can deal with the concept of a hierarchy of perception and control by using more or less concrete examples. I still don't know how I could have done it differently. Well, of course I could have guessed right, I suppose -- but what if I haven't?

That's what we both have to ask, and the only way to answer that question is to get out the picks and shovels and start doing some basic science.

How is behaviorism different from PCT, and how can we find out which one is the better description? To ask those questions is to take a great risk, because it implies that we will submit to the answer. The answer could go against you or against me. That does make me nervous, I admit. If you have the brains I think you have, it will make you nervous, too.

On the other hand, if you and I and those who work with us could come to an actual agreement, the result would be of immense importance. The result would be far more significant than two guys and their devoted supporters playing 'tis so, 'taint so. It would be a triumph for science itself. I would love to be part of that. My naive highest ideals -- OK, values in your terms -- would be justified and vindicated in the face of cynics. I can even see that the result would be freedom from the bonds of short-sighted self-interest, which I might deny under other circumstances, but which I can feel.

Before I ask where we start, I'd better check with you about the basic proposal. Is this worth doing?

Best,

Bill

I'm off to Sunday brunch with the family so time only for a brief response

One reason I think PCT and ACT arrives at much the same clinical point is that a) we seem to be building these systems in a similar way -- i.e. bottom up and process focused, b) behaviorism in the hands of functional contextualists share some similar analytic units.

My views on a) seem solid as I look at what you do; b) is a little more iffy. b) is more iffy because it seems that PCT is a mechanistic theory (or an elemental realist theory if the word "mechanism" seems pejorative, which it is really not meant to be in this context) and so much of the ACT work is consciously contextualistic and ontological. That is a long story. Stephen Pepper's system helps make sense of it. I've attached an article on him and if you want to examine it.

So I suspect we might differ at the level of philosophy
But at the level of actual units, I'm not so sure. It seems similar at least underneath the surface.

The words get in the way of course.

Take the idea of a 4 term operant. It's a unit. 1 whole unit. you can't fully disassemble it and still have the "it." You construct the unit for pragmatic purposes so it is not an ontological "it" -- its just another example of the *scientist* doing what you say people are doing (that is the key insight of so called "radical" behaviorism -- but it overturns most of traditional behaviorism to do that so radical contextualism might have been clearer)

The concept of "behavior" in a 4 term operant changes from a topographical descriptive ontological thing to a class based, functional aspect of a large unit.

Let's take a classic behavioral concept like generalized imitation viewed as an operant.

Generalize imitation is something like a learned set of functional relations:

Given-that-there-is-a-need-and-purpose-doing-what-other-people-do-based-on-its-utility-in-accomplishing-such-purposes

The hyphens are in there on purpose.

There is a Establish operation / discriminative stimulus / functional action / regulating consequence relation nested inside the hyphenated clauses and learning to transition among these view as a learned skill: "the operant." It's one whole unit.

Is this a control system? Seems rather like one.

We can break it down into elements for pragmatic purposes but the elements are no longer sensible in the absence of the overall system. For example, there is no such thing as a discriminative stimulus in the absence of the other three terms; the same with the others *including these classes of functional actions*. Since behavior without the system cannot be broken down into sensible units one has very little to say about it. It's not even an "it." It's just part of the one; an unanalyzed aspect of the whole.

That kind of behaviorism is rather different than the mainstream of behaviorism people think they are talking about when they say the "b" word

One reason I call this wing "functional contextualists"

The idea of nested levels fits with in this in a broad perspective (without the clear delineation of levels as in PCT). Skinner talked about writing a book as an operant; driving to the beach as an operant; etc. But each aspect could be broken down into the same systemic analysis. For example, if we are looking at an action like driving to the beach, we can become more interested in how the person steers or how they saw and reached for the door handle. In each case the same case hyphenated functional analysis is applied. Skinner has this cool little piece on how seeing depth cues is learned in the early period of an animals life. It is nested inside selectivist views of biological structure -- e.g. we are prepared to see depth cues; some of that is not learned etc but on top of that we literally learn to see

So back to my main point

If I am right that (despite what appears to be some philosophical differences) PCT and our weird wing of behavioral thinking is fairly compatible, then the similarities in MOL and ACT might be accounted for by the process focused, bottom up strategy overlap; and the relentless focus on nested control systems / functional system ("operants" the way I'm talking about it).

Why is ACT different than other functional behavioral systems? It's mostly because we did three key things: we refused to trivialize complexity (we thought that happened in early behavior therapy); we agreed that language and cognition had no adequate functional account; and we developed such an account.

One great advantage of functional contextual thinking is that we can easily avoid "'tis so; tain't so" because we view ourselves as just playing a game. We are balancing the stick on our hands. We are jumping that bicycle onto the fence. That is made easier by abandoning truth with a capital T; or ontology with a capital O. And we've done that -- so we are already naked, in a sense.

"Freedom is another word for nothing left to lose"

Once you lose what is REAL things seems a bit more free (we are realists, but only in the limited sense of assuming the one world ... which you might as well call "real").

So whatever works, works.

My biggest worry re PCT is the limited amount of outcome data in the applied realm. I hope that does not sound offensive. It is just that there are so many ways to get it "wrong" in the more pragmatic sense of "did the stick fall or not" or "did the bicycle hit the ground or not."

So I'm left more with how this kind of play seems to help me balance that bicycle ... or in our more usual way of saying it how it helps me "predict and influence the whole organisms interacting in a with a context considered historically and situationally using principles that have precision, scope, and depth" ... or even more grandiosely how it helps me "create a comprehensive psychology more adequate to the challenge of the human condition." That phrase is the adopted valued goal of ACBS

I do have a sense that PCT adds something.

So I'm willing to play.

I'm also buried with demands and slow, but that is OK. There is not a speedometer on our foreheads. Except the one that says "if I don't get ready for brunch my wife will kill me"

I'm out

Peace, love, and life

Pepper review JEAB 1988.pdf

<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1338844/>

Hi, Steven --

The review of Pepper is both interesting and hard for me to grasp. I see a link between his "world view" and what I call a "system concept." All the properties of world views appear to apply to system concepts, including the fact that (in different people) they tend to be self-contained and unconnected to each other, though of course they are built on lower-level experiences that may be similar. When different and mutually-exclusive system concepts appear inside an individual, the person is probably in some sort of deep trouble.

The four types of "relatively adequate" world views are not exhaustive or mutually exclusive, as evidenced to me in the fact that the theory embodied in PCT has clear features of all four of them.

PCT includes mechanisms, forms, organicism, and acts in context. An act in context is hinted at in the PCT saying, "You can't tell what a person is doing just by looking at what he is doing." A person waving an arm up and down may be cooling a burnt finger, signalling hello to a friend, conducting a symphony orchestra, hailing a taxi, voting, or sending a semaphore message. PCT is organic in that a control process is carried out by a set of functions, no one of which can be said to control anything at the level under examination, although one of them (the action generator) may be composed of lower-order control systems. It embraces form in that a given controlled variable can be a function of a small or large collection of distinct lower-order perceptions, and at the same time itself be one element among many of a higher-order perception -- perceptions being the variables that we control. The visual doughnutness of a perceived object may be seen in any doughnut large or small, frosted or plain, and to varying lesser degrees in what we describe in other contexts as an inner-tube, the letter "O", a smoke-ring, and a quoit. Many different perceptions, whether or not they have any shared characteristics, come to be perceived as members of the same class, to which we give names like "doughnut" or "gorilla" or "contents of my pocket" (using the visual or auditory configuration-perceptions we call "words" as the symbols representing other kinds of perceptions). This aspect of form relates in PCT to levels of perception. Currently, there are 11 kinds of "forms" in PCT, only one kind, configurations, including what we would perceive as a physical object.

And of course PCT includes neural mechanisms, most still only a short way beyond the proposal stage, which explain how it is that any of these kinds of systems can physically exist and function in a real organism as we observe them to do.

I think that the levels of perception in PCT take the same whole that is described by Pepper's world views and separates its components along planes of cleavage that are more natural than Pepper's components. That is because PCT does not consider the elements of these four categories as distinct from one another, but simply as different aspects of the whole system. Starting with the whole system, PCT describes levels of control of many different classes, each level including a multitude of semi-independent control systems with each controller composed of functions we identify as perception, comparison (evaluation), and action. The "action" component consists of processes for specifying the reference conditions or goals to which lower systems then try to match the perceptions that they control, by manipulating the goals for lower systems still, or at the bottom, muscles and glands. All the levels act simultaneously, so to grasp how the system works it is necessary to consider it as an organic whole. Yet PCT models are built from the bottom up, and show how it is that lower-level processes can yield consequences that carry out higher-level intentions.

In short, PCT would seem to encompass most of the philosophical and practical considerations you have brought up. For that reason, I am hopeful that you will find in it a useful tool for your purposes.

There is, however, a price attached to understanding and using PCT. If PCT is the right approach to understanding organisms, then all the other approaches are not just different; they are wrong.

The amount of wrongness varies from a few percent to 100%. Every well-educated person in the world who now supports PCT as the theory of choice started out believing in some other idea of how the behavior of organisms works -- at least half of them to the point of having a PhD based on some other understanding. Every one of them had to make decision after decision where there were conflicts between old ideas and the PCT view. A very few of them (that I know of) decided against PCT and left. But nearly all of them are still here, 50-odd years after the first publication.

You're not going to be an exception, I'll wager, though I won't try to predict your decisions. Neither do I want to be the one who criticizes any idea you support. Only your criticisms matter.

You say that the truth criterion of contextualism is successful working. That's my criterion, too, for believing in any theory. A theory purports to explain behavior. Very well, let's use the theory to accomplish the particular goal of creating predictions of what will happen when something else happens, and see if the theory works. I am very big on demonstrations, as you have probably figured out by now. Most of my demos involve a human being doing something, so the outcomes are not simply calculations that obey whatever rules I choose to use. If we disagree about something, I (and Rick Marken and Bruce Abbott and Yu Li) can probably devise an experiment to test which point of view works. I'm pretty serious about this: if my idea doesn't work, out it goes. If your idea doesn't work I expect you to say the same thing, but I'll try very hard not to say it for you if you don't. You can trust me to honor that promise for an inch, but not for a mile.

Unless you see some more fundamental place to start, I suggest that we begin with talking about "the operant." Or perhaps it should be the "reactive organism model."

Best,

Bill

Hi, Steve --

I guess I'm anxious to get started.

SH: So I suspect we might differ at the level of philosophy

But at the level of actual units, I'm not so sure. It seems similar at least underneath the surface.

BP: Same here. However, the two areas overlap. There are kinds of real systems made of matter and energy that literally can have purposes of their own, select stimuli, and do all those other bad mentalistic things that the concept of the "reactive organism model" rules out. No "scientific" psychologist from Skinner on back, of course, understood that such things could exist -- they thought they were mystical fictions or had religious connotations.

SH: Take the idea of a 4 term operant. It's a unit. 1 whole unit.
you can't fully dissemble it and still have the "it."

...

There is a Establish operation / discriminative stimulus / functional action / regulating consequence relation nested inside the hyphenated clauses
and learning to transition among these view as a learned skill: "the operant." It's one whole unit.

Is this a control system? Seems rather like one.

BP: We can model it that way, but we still need ways to find out if the model is correct. In a control model, the main components (in terms of what we could observe from outside the behaving system) are a controlled variable, a reference condition, an action, and a disturbance (an

independent variable affecting the state of the controlled variable). No one of these "controls" at the same level of description, so this fits your idea of the operant. You can't disassemble these components and still have a control system. Even the disturbance is undefined without any specification of the controlled variable and the reference condition, both of which are observable from outside the behaving system.

In PCT we don't simply assume that a behavior is part of a control process. We have a "test for the controlled variable" which can be applied with varying degrees of rigor to see whether the criteria that identify a control system are present. My interactive models all use an advanced form of this test.

The simplest type of test is simply to apply a disturbance to some variable in the environment of an organism, a disturbing variable selected so we could predict the effect it would have on the variable we think might be controlled if there were no control system controlling it. If we observe exactly that effect, we have shown that this is not a controlled variable. If the predicted effect is completely missing, we have a perfect control system -- somewhere.

Example: I slowly push a vase on a table toward the edge of the table. I predict that this disturbance of the position of the vase will, if nothing is controlling that position, cause the vase to topple off the edge of the table and fall to the floor. As soon as the vase does that, I can say there is (was) nothing presently controlling the position of the vase.

On the other hand, since I'm actually doing this to test some kind of hypothesis about an organism, not a vase, I have my eye on the person who owns the vase, and if that person is controlling the position of the vase I would expect an intervention of some kind before the vase actually topples. I would therefore not be surprised to hear this person say "Hey, watch out, what are you doing?" If I don't pay attention to that, but keep slowly sliding the vase toward the edge (with my elbow so it looks accidental), I expect the action from the person to become stronger and more effective: a hand reaches out and snatches the vase away from danger. At that point I can apologize for my carelessness, make up some lame excuse, and walk away with some degree of confidence that I have been disturbing a controlled variable and that the person who acted was doing the controlling. The action produced by the person did what a controlling action should do: it kept the disturbance from having the effect on the vase I predicted from knowledge of the physical properties of the environment and the assumption of no control.

Of course that is only a first approximation. I would need a new subject for each subsequent refinement, because no subject would let me get near a vase a second time, but the next time I could watch to see exactly where relative to the edge of the table the vase was when the controlling action began, and in further experiments, how the strength of the action varied as the distance to the edge decreased. Also, I could test a prediction based on the fact that a control system has to sense the variable that is being controlled; I could arrange matters so the person I was investigating couldn't see where the vase was; that ought to prevent control, and of course it most likely would, though I would have to be prepared to catch the vase when it toppled to avoid having to pay for it out of a nonexistent research budget. If the hand did snatch the vase away when I was preventing his perception of the vase, I would have to revise my idea of what is going on -- perhaps he sees the vase in a mirror, or its shadow, or perhaps he sees his wife (who can see the vase) frantically signalling him to do something. Maybe it was his wife's hand I saw. If I can't find out how he manages to react appropriately without seeing the vase, I just have to scrap the whole model and start over. But the chances are pretty good that the experiment would actually work.

By repeated experimentation I could find a population average for the nature of the controlled variable, the reference condition defining the limiting value of the variable at which the action is still just zero, the action that occurs when the actual position goes past the reference condition, and the quantitative relationship between the action and the disturbance. That would give me enough material to put together a computer simulation with a standard control system in it, and adjust the parameters of the simulation so the model's behavior came as close as possible to the real

behavior. This experiment isn't conducive to accurate modeling, first because of having to use a different subject for every trial (the fit would be terrible because of using population measures), and second because simulating complex behaviors like reaching out and grabbing the vase would tax my ability to analyze physical systems.

The same principles can be investigated in simpler nondestructive settings that allow doing in-depth investigations with a single subject, which is why most of my work uses computer simulations and control of screen displays.

So that's how I can say I think I know that some particular behavior is a control behavior, and speak about sensing, acting, and resisting disturbances. The PCT model itself is simply a proposed organization of nervous-system and motor actuator functions that can accomplish the observed kinds of control.

"Control," by the way, has a very specific technical meaning in PCT which has to be distinguished from its various common-language usages. If A affects B, and if an increasing disturbance of the state of B results in an action by A on B that prevents most of the effect of the disturbing variable on B, then A is said to control B. If this is accomplished without A's having to sense the cause of a change in B (so A's action is based strictly on sensing the state of B), this is said to be classical negative feedback control, the kind we use in the PCT model.

I would venture that the establishing condition (which tends to alter the controlled variable) corresponds to the disturbance, the functional action to the effect of the behavior on the controlled variable, and the regulating consequence to the state of the controlled variable (which in PCT does no regulating but is itself regulated by the control system, as is easy to demonstrate). I'm not sure how to handle the discriminative stimulus -- perhaps it would be a higher-level controlled variable or what you might call a "contextual" variable.

SH: We can break it down into elements for pragmatic purposes but the elements are no longer sensible in the absence of the overall system. For example, there is no such thing as a discriminative stimulus in the absence of the other three terms; the same with the others *including these classes of functional actions*. Since behavior without the system cannot be broken down into sensible units one has very little to say about it. It's not even an "it." It's just part of the one; an unanalyzed aspect of the whole.

BP: That may have been the case prior to PCT, but it's not true any more. I think we have identified a set of elements that, together, adequately explain the behavior of the whole. A model composed of these elements and their properties can, once set in motion, very closely approximate the behavior of a real person. Only in fairly simple examples, so far, but we're learning how to handle more complex ones and see no reason to say we won't ever do it. My next Holy Grail to seek is a model of a whole body that can roll over, push itself into a squat, stand up, walk, and run.

SH: The idea of nested levels fits with in this in a broad perspective

(without the clear delineation of levels as in PCT).

BP: The idea of levels is more important than the specific ones I've proposed.

SH; Skinner talked about writing a book as an operant; driving to the beach as an operant; etc.

BP: Each of these is what we would call a controlled variable: a perceivable entity that is being brought by behavior to a specified condition.

SH: But each aspect could be broken down into the same systemic analysis.

For example, if we are looking at an action like driving to the beach, we can become more

interested in how the person steers or how they saw and reached for the door handle. In each case the same case hyphenated functional analysis is applied.

BP: The behavior that accomplishes steering or grasping door handles is essentially unpredictable: it depends almost entirely on what random disturbing variables happen to be present in the environment. It is not necessary for a control system to sense those disturbing variables; all it needs to sense is the controlled variable itself. What is reproduced is not the behavior, but a consequence of the behavior. I can demonstrate many situations in which the joint consequence of actions and disturbances is accurately repeated, even though the disturbances are randomly generated and the actions do not correlate with the controlled variable. This situation, highly non-intuitive in most theories of behavior, can be demonstrated beyond doubt.

Skinner sensed this fact about behavior but didn't see the correct explanation (and in fact rejected it when he encountered control theory). Instead, he proposed that behavior is controlled by its consequences (which is how it looks), defying the simple observation that it is the (variable) behavior that controls the (repeatable) consequences. I think he suspected the existence of the feedback loop but couldn't quite grasp how it worked. He picked a solution that fit the prevalent idea that organisms, being made of matter, obeyed only the laws of physics and chemistry, so causation had to run from the environment, into the organism, and out again into behavior. Then behavior would alter the environment and another cycle would start.

=====
That's a big enough chunk to chew on for now,. There's more of interest in your post but it will come up again. I hope I haven't already broken my promise to let you do the criticizing of the ideas you present.

Best,

Bill

Sorry. I've been way busy with end of semester stuff. I am starting my sabbatical and will be in Europe most of the summer beginning in two weeks, with multiple workshops and the like so I've been preparing like mad and haven't been able to really drill down in Bill's message.

I'm not sure how many installments will follow until I have time to check out the MOL work. That may take a while. As I said earlier in the exchange though, it is a big barrier to me that there is not an extensive empirical literature on MOL because that is where the rubber meets the road for an applied person like me.

There are many interesting overlaps that I see -- there are also deep differences in assumptions and models of truth.

When Bill says

"That may have been the case prior to PCT, but it's not true any more. I think we have identified a set of elements that, together, adequately explain the behavior of the whole" what he is saying is that he has a mechanistic model that works.

I think PCT has as much or more of a chance of that than any mechanistic model I've yet examined but a) "works" has to be evaluated at the right level.

The kind of working I want I do not yet see anywhere in the PCT work (e.g., see above).

It is more where you are headed than where you are, and b) data cannot do everything. As Godel proved, any symbolic system has to start outside the system with assumptions. And there the differences are large.

Bill earlier pointed to the cool possibility that we could reach over boundaries and come to agreement. It seems to be more likely that we could reach a point of

mutual appreciation. PCL and CBS ("contextual behavioral science" -- the name for our wing of the behavior analytic tradition) are different in their basic assumptions.

If I say "the whole is primary; the parts are derived" that is a philosophical assumption. If Bill say "no I can model the whole from the elements .. the elements are primary" that is an assumption. There aren't data that can address assumptions of that kind ...

Anyway ... I apologize for the delay. More when I get examine MOL

Hi, Steve --

SH: There are many interesting overlaps that I see -- there are also deep differences in assumptions and models of truth.

BP: That's OK. I assume you have reasons for both the assumptions and the models, so we don't have to just disagree by yelling louder, with the loser being the first guy who runs out of breath (that would be me -- you saw the oxygen bottle). Tell me what you assume and why, and I'll reciprocate. No reason why we couldn't work out the differences.

SH: When Bill says

"That may have been the case prior to PCT, but it's not true any more. I think we have identified a set of elements that, together, adequately explain the behavior of the whole" what he is saying is that he has a mechanistic model that works.

BP: Actually, it's a functional model when looked at from the top down, mechanistic as seen from the bottom up. You really need both to have a good theory.

From the top down, we start with an abstract mode of perception such as perception of principles, and a reference condition (akin to your "values") which is the desired or intended state selected as a goal for that perception. You want to be honest, but just the right amount of it, not *too* honest. The difference between perception and reference is what drives changes in more concrete goals as the means of altering the higher-level perception toward its given reference condition. That's the basic principle: alter the action and its effects on the environment to make the perception of the effects more like the desired perception. It's also a hierarchical principle: the means of altering the perception is not to produce physical action directly, but to alter lower-level goals which relate to the means of altering the higher-level perception. Only at the lowest level is the action an actual muscle force.

From the bottom up viewpoint, the problem is to find a reasonable design for the kind of physical system that can behave like this. What I have proposed as the basic kind of physical system is a negative feedback control system, first analyzed and understood in the 1930s by electrical engineers. It seems to provide exactly what is needed -- so far.

I have working models that show the feasibility of this approach for modeling simple behaviors like manual control or visual-motor tracking of targets. Others, like Robertson and Goldstein in our "Control Systems Group," have done proof-of-principle experiments at higher levels -- e. g., a simple experiment predicting behavior when an expressed self-concept is contradicted by the experimenter.

From the top-down viewpoint, there is a good deal of Applied Control Theory under way (how do you like that? "ACT"). Warren Mansell's group in England, Tim Carey in Australia, a number of people including David Goldstein in the USA, are all using principles of PCT and the therapeutic approach of MOL. We're not organized as well as you and your group are, but we're slowly getting there.

SH: I think PC[T] has as much or more of a chance of that than any mechanistic model I've yet examined but a) "works" has to be evaluated at the right level. The kind of working I want I do not yet see anywhere in the PCT work (e.g., see above). It is more where you are headed than where you are, and b) data cannot do everything. As Godel proved, any symbolic system has to start outside the system with assumptions. And there the differences are large.

BP: Maybe not as large as they look right now.

SH: Bill earlier pointed to the cool possibility that we could reach over boundaries and come to agreement. It seems to be more likely that we could reach a point of mutual appreciation. PC[T] and CBS ("contextual behavioral science" -- the name for our wing of the behavior analytic tradition) are different in their basic assumptions.

BP: It's the behavior-analytic tradition that we will have to work on the hardest. Let us remain cool. From the PCT view, it's a plausible misinterpretation. I think I can demonstrate and support with evidence a more useful interpretation of the data.

SH: If I say "the whole is primary; the parts are derived" that is a philosophical assumption.

BP: It's also a debatable statement of fact. From the top down, the whole is primary, in that the higher-level control processes determine what the lower levels are going to try to accomplish. But when you ask how it is done, you also have to consider the bottom-up approach.

The higher levels in PCT are emergent from the lower levels but are not reducible to what the lower levels do. I think you will agree with that principle, and I can demonstrate with tangible models exactly what that means. On the other hand, there is no magic; effects can't reach back through time to alter their causes, nor does control theory required that. However the system works, it doesn't violate any laws of physics, at least the ones we know of and have stated correctly. When Skinner said that behavior is controlled by its consequences, he was simply wrong. Consequences, the ones that matter most to the organism, are controlled by behavior, in the exact technical sense of "control". Go ahead and say "But... but ..." but you will agree with me in the end.

SH: If Bill says "no I can model the whole from the elements .. the elements are primary" that is an assumption. There aren't data that can address assumptions of that kind ...

BP: Yes, there are. Of course as the data become clear, the definitions will morph into new definitions, and that is how we will ultimately achieve agreement. The meaning of "primary" will change.

My goal is not just to develop a theory of behavior. It's to develop a theory of behavior that is consistent with the best of all our other theories about real phenomena and all the data we know about. Psychology can't afford to ignore physics or neurology and vice versa all three ways. Our understanding of one part of nature can't be in contradiction of our understanding of any other part if we want correct theories. That's a basic principle, to me. But I don't put any discipline at the top: they all have to adjust to each other to achieve true understanding.

Good luck with your sabbatical and don't forget to enjoy yourself in the middle of all that work.

Best,

Bill
