BIG DATA, AI, AND MACHINE (SELECTIONIST) LEARNING: A STROLL THROUGH THE THINKING OF ADVANCED PROGRAM DESIGN

T. V. Joe Layng, Ph.D.

I must admit some reluctance when I was approached about writing a blog post for the Technology SIG. My first reaction was that I had little to contribute beyond what has been already written on the blog and elsewhere. In education, there are many useful articles that distinguish between technologies of tools and technologies of process, including games and gamification. In clinical work, the use of bug in the ear, video, and a variety of sensors have been explored. And an increasing number of apps help users set goals, monitor progress and communicate with health or behavioral professionals.

However, as I spent time thinking about all this, I began to think about a topic that might have some value for the readers of this blog. What if—instead of an entry about how technology could facilitate behavioral practices (or vice versa)—I described how I would approach a development project that required an outcome not possible without the intertwining of behavior analysis and certain technologies?

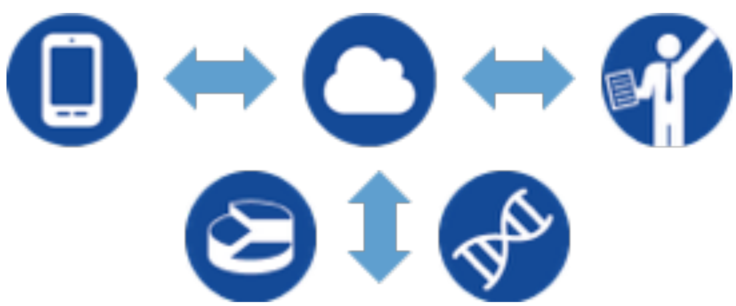

The Project

The project I will write about was proposed and planned but never completed, so we will never really know if the approach would work. As with every development project, the data generated from testing with users are what actually shape the course of development, though a plan is usually developed that at least guides initial efforts. It is this plan and how user testing influences design and development that will be the subject of this series. Although the project never moved passed the planning and initial design stage, the thinking and approach described is real in the sense that it is what would have been done.

Today there are countless apps designed to help people meet individual goals. Apps provide advice, offer easy interfaces for collecting personal data, include clever graphics to represent progress toward a goal, provide feedback upon subgoal completion, award points, badges, or other acknowledgements, and may provide social support by linking to others. Some apps are designed to help individuals collect data and communicate with therapists or others who are counseling or coaching them. In almost all of these applications, the app itself is a conduit for gathering or evaluating data and providing feedback. In short, some behavior analytic principles may be incorporated into the app itself, but often the app collects data for a therapist or counselor who is implementing interventions that would likely have occurred if the same information had been gathered in another way. In essence, technology is used to provide faster—and perhaps better—data that are analyzed and acted upon much the same way as they would be without the technology.

Would it provide needed services? It likely would. But, nothing would be invented, nothing would really change the way service is delivered, and no tools would be created that might inform and improve behavior analysis itself.

I was part of a group asked if there was a way to affect large scale health-related behavior change using technology. The first impulse was to do exactly what was described above—design apps and digital job aids that would allow fewer counselors to help many more people. This would include remote interviewing, distance coaching, digital call centers, and apps to record data and provide feedback. Technology would play the role of extending human reach. In truth, this approach, though not without its challenges, is a pretty straightforward application of behavior analysis and existing digital technology. Would it provide needed services? It likely would. But, nothing would be invented, nothing would really change the way service is delivered, and no tools would be created that might inform and improve behavior analysis itself.

Discontent with the Status Quo

I began thinking in earnest that we needed to do something entirely different. Could we use the emerging areas of machine learning, big data, and selectionist algorithms to actually enhance our procedures? Could we actually go beyond simple linear behavior analysis and bring a more complex nonlinear analysis and systemic interventions (Layng, et. al., 2022) to large numbers of people, without having to extensively train therapists, counselors, or coaches? In a 1968 paper, behavior analyst Israel Goldiamond and psychoanalyst Jarl Dyrud speculated that one day perhaps we could develop what they called programed clinical instruction (PCI). PCI would consist of programs targeting repertoires the absence of which was the reason for seeking therapy. Could their 1968 speculation be turned into reality with today's technology? Several new applications of existing technology along with the invention of new technology combinations would be required.

Big Data

The first step was to think about where we needed to go. What was it that we wanted to see at the end of development? As I investigated big data applications, I found they were providing some fascinating results in many fields. For example, I spoke with the CEO of a major marketing firm about their application of analytic algorithms to big data. In big data applications, seldom is there a prejudgment about what data are important. The more data the better. They use many sources—local traffic information, zip code demographics, census data, crime rates, pet ownership, and on and on. In this case, they used big data to identify the person most likely to use the facial cream product they were hired to market. They found that person to be a middle-aged woman with an upper middle class income who lived in the suburbs and who owned a dog. Move the woman to the city, the likelihood dropped off. Take away the dog, the likelihood dropped off. Why was this the case? No one knew and they didn't care. They targeted this segment with advertising and promotions directed at this "person" and sales took off.

We soon found there were similar applications in healthcare. Hospitals in a northeastern state were having too many readmissions to the hospital within 30 days from discharge. With the help of a big data firm, they determined that older women from their area who are discharged to the home of a lower-income, working, grown child where the spouse is also working have an 80% chance of being readmitted to the hospital in 30 days. What was fascinating was that this was true regardless of admission diagnosis. If only one of the couple was working or there was no spouse, if the patient was discharged to their own home, had a slightly higher income, or was slightly younger, etc., the readmission rate was much lower. Again, one may speculate as to why this was the case, but for the hospital it didn't matter. They began home visits for this group as well as follow-up telephone communication and greatly reduced readmission rates as a result.

Genetic Algorithms

It was clear that this type of big data analysis could be very important in identifying what procedures might work with what individuals. Equally interesting was that the hospital data were generated by algorithms that had not been directly programed. They were genetic algorithms that evolved based upon feedback from their increasing predictive success. Now that was exciting. Gathering large amounts of seemingly disconnected data, evaluating client success through the program, and feeding this back into evolving genetic algorithms might lead to being able to better match people to interventions. We could look at existing data on people who did or did not succeed in various life improvement programs and see if we could find initial big data correlations that appeared to predict that outcome. Perhaps we would not have to start from scratch.

Equally interesting was that the hospital data were generated by algorithms that had not been directly programed. They were genetic algorithms that evolved based upon feedback from their increasing predictive success.

Application of big data analytics utilizing genetic algorithms became part of the vision. Could we, for example, pre-identify people who would fall into the following categories:

- Self-Motivated: give them information and they run with it.
- Motivationally Challenged: people who express interest in change, but can't seem to do it on their own.
- Clinical Psychological Issues: people with serious life issues besides the targeted health-related patterns.
- Fatalistic or Resigned: those who believe nothing will work.

By identifying individuals falling into each category, initial individual assessments and programs could be customized. For example, for someone falling into the self-motivated category, simply providing initial information, digital guidelines, and rudimentary data collection and feedback may be the only intervention required. For the motivationally challenged, a different assessment and program would likely be required. I will return to this discussion later in the series.

The Constructional Questionnaire

Next, we would need to individually assess the people participating in our behavior change programs. Now there are all kinds of surveys and questionnaires that are used by psychologists, behavioral or otherwise. Often, these questionnaires are trying to operationalize some form of internal hypothetical construct, such as resilience or psychological flexibility. While perhaps useful under some circumstances, this form of methodological behaviorism–an operational definition based on an observable indicator response representing an internal hypothetical construct–would not be adequate to our task. We needed an assessment that helped elucidate the sets of alternative consequential contingencies that made sense of the disturbing or problematic patterns our clients were displaying. Israel Goldiamond had developed just this sort of questionnaire. It was called the Constructional Questionnaire. It is a powerful tool that can help identify the alternative sets of contingencies responsible for the current patterns as well as provide an evaluation of the client's current relevant problem-solving repertoires. But, not unlike many behavioral procedures, it requires a trained nonlinear behavior analyst to effectively use it. Accordingly, it seemed that the analysis provided by the Constructional Questionnaire could not be delivered at anywhere near the scale required for our project. However, in rereading the case presentation guide provided to trainees, I was struck by an idea about how it might be possible after all.

The guide began with, "Weave in various items from questionnaire and other sources to present a coherent picture of a person functioning highly competently, given his circumstances and implicit or explicit goals." The reference to "coherent picture" immediately jumped out at me. The extensive reading about genetic algorithms and selectionist approaches to big data analysis combined with seeing the word "picture" worked together to produce an exciting possible solution as to how a Constructional

Questionnaire could be used at scale. In essence, the next element of our emerging vision might indeed be possible.

What precisely was the solution? How did it occur to me? We continue our stroll in Part 2 of this series.

PART 2

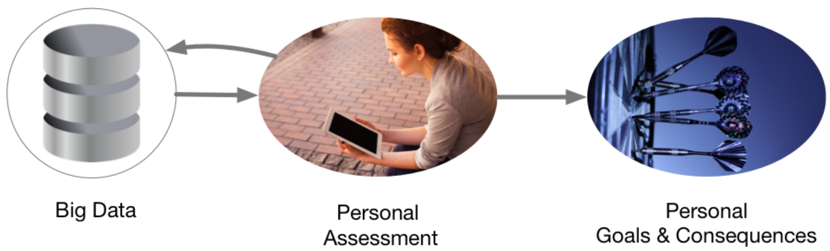

Weave in various items from questionnaire and other sources to present a coherent picture of a person functioning highly competently…

That word—picture—combined with the context of selection-based algorithms brought responding to a book I had read some years earlier to strength.

Okay, it made me think of the book Why We Feel: The Science of Human Emotions by Victor Johnston.

In the book, Johnston describes the use of selection-based algorithms to produce a likeness of a person only briefly seen. Producing such a likeness is routinely performed in police investigations and is typically done by an observer describing facial characteristics to a graphic artist. Instead of this typical method, Johnston displayed on a computer screen various combinations of facial features. The observer then rated the characteristics and the face as to how well they matched the actual person observed. The algorithms generated variations, selected against variations that received low ratings, and presented new combinations based upon known variation principles. Within a short time, the algorithms provided a picture of the person that was a very close match to the person observed and was, in fact, much better than the artist's sketch.

The immediate question was raised: could a similar procedure be used to provide "a coherent picture of a person functioning highly competently…?" Of course, a questionnaire is based on asking questions, but so is an artist's sketch. Could a series of multiple choice or yes-no questions be used to generate an early picture—a written scenario—that could then have its features and overall description rated? As a result of interacting with the program, could a picture—much like what is derived from the Constructional Questionnaire—be accurately produced (as rated by the user)?

The Constructional Questionnaire begins:

“I am going to ask you a group of questions about your goals. You are here because you want certain changes to occur, or want something else. (a. Presented outcome) The first of these is: Assuming we were successful, what would the outcome be for you?”

At this point the client typically talks in generalities about being less anxious, depressed, etc. Or if more positive, “I would be a better communicator.” The client is allowed to say whatever they like with little direction.

The next series brings much more specificity to the answers by asking what precisely the client would be doing if they were not anxious, depressed, etc. This is not always easy. People are used to speaking pathologically. That is, they can easily say what they would like to stop doing. But, describing what they would be doing if the present problem were not a problem can be challenging for many. To help with this, the next question was creatively designed to produce useful answers:

“Now, this may sound silly, but suppose one of these flying saucers is for real. It lands and 2,000 little Martians pour out. One of them is assigned to observe you—your name was chosen by their computer on some random basis. He lands some time after L-Day —Liberation Day from your problems—and follows you around invisibly. He records his observations and these are sent back to Mars. Their computer will decide on the basis of the sample of 2,000 Earthlings they have what their disposition toward Earth should be. What does he observe? (Alternate or added form: What would others observe when the successful outcome was obtained?)”

The client is then asked to begin in the morning and describe what the Martian sees from awaking in the morning to falling asleep at night, on weekdays and the weekend. Remember, this is life after the problems are solved. When a patient says, “I would have a happy exchange with the receptionist,” the interviewer responds by saying, “What does the Martian see when you are having a happy exchange?” Eventually, a picture of how the client’s life would change is provided. It is comprehensive and thorough and is likely something to which the client has given very little thought.

The next question asks, “How does this differ from the present state of affairs? Can you give me an example?” Here is where the client describes and gives examples of the difference between the way things are now and where they would be if all was well. It provides the before and after from which the success of the intervention will be determined.

Other questions are designed to explore other important areas, such as whether there is a “hidden agenda” not mentioned in the outcome question, what is going well and would not change, and one which often provides revealing answers: “You’ve heard of the proverb, ‘It is an ill-wind that blows no good.’ With regard to some advantages that might have ‘blown your way,’ has your problem ever produced any special advantages or considerations for you? (Examples: in school, job, at home) Please give specific examples.” The full questionnaire may be seen here: http://journals.uic.edu/ojs/index.php/bsi/article/view/92/117.

Here was a method that had the potential to bring the benefits of this very sophisticated interview to thousands of people.

Two possibilities immediately emerged. First, could we ask some carefully constructed multiple choice type questions that would allow us to construct a completed scenario for the question about successful outcomes, which could then have its components as well as its overall picture rated? We could iterate this process using composite scenarios with the selectionist algorithms producing variants until the components and composite received high ratings for accuracy.

The second approach would be to take smaller steps. Instead of constructing a composite scenario and rating it’s various components and the overall picture presented, this approach would present a component for rating. Once the component evolved into a highly rated one, another component would be added and the process repeated until a composite was produced. For example, “What would the Martian see when you awake, arrive at work, at lunch, etc.?” Mini scenarios would be generated and put through the process. Responses to earlier questions may be useful for producing subsequent variants. These components would eventually be put into a larger composite for evaluation.

Would one approach or the other yield better results? Would both yield the same result, but with one more efficient than the other? Could it be done at all? As to the last question, given my investigation into how Johnston contracted his algorithms, I was confident it actually would work. The accuracy of the scenarios could be evaluated using ratings as well as by comparing the results to verbal pictures produced during actual interviews. Once there was some confidence in the outcome, an evaluation could be conducted to determine how well programs whose outcomes were defined by this automated selectionist interview actually performed when used in a program to help individuals meet their goals.

Think about it for a minute. Here was a method that had the potential to bring the benefits of this very sophisticated interview to thousands of people. An interview that specifies not only the consequences maintaining the client’s disturbing patterns given their alternatives, but that could specify outcomes having the same or better benefits at less cost to the individual. Further, it might be possible to produce different versions of the interview to match individual characteristics determined by a big data analysis.

Earlier I described four possible categories into which potential uses might fall. Each person in a particular category could be provided with a separate interview customized for their category. Our research and development plan would need to coordinate these two features. Further, our individual assessment results would be fed back to the big data algorithms to help define and make more precise the categories and their predictive validity.

This would not simply be an app that provided goals and feedback, but one that evaluated progress, collected the right kind of information (including client emotions), adapted to user performance, benefitted from big data analytics, and would be able to switch from "topical" to "systemic" approaches to achieving client outcomes.

So now two elements of what would be required were specified: the big data feature and the selection-based automated interview based on the Constructional Questionnaire. But once the contingency analysis was made and a case guide such as described in Part 1 was generated, a delivery program would be required. This would not simply be an app that provided goals and feedback, but one that evaluated progress, collected the right kind of information (including client emotions), adapted to user performance, benefitted from big data analytics, and would be able to switch from "topical" to "systemic" approaches to achieving client outcomes.

Now an app would have to be created that could provide both topical and systemic programs, collect personal data, and connect everything to big data analytics and ongoing selection—the topic of Part 3.

PART 3

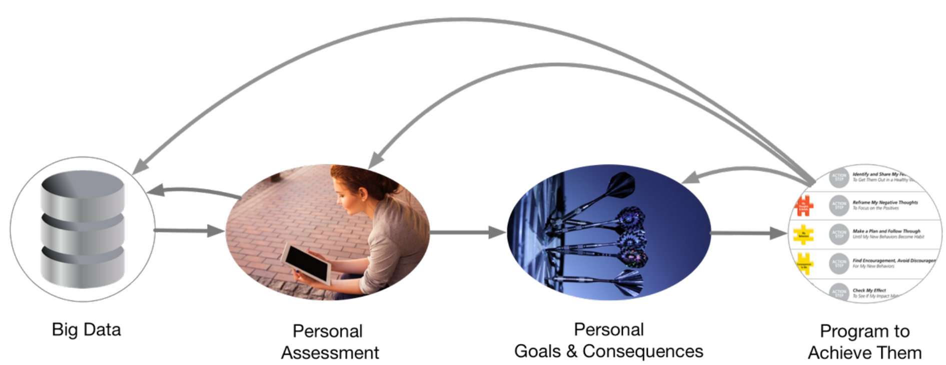

In Part 1 of this series, I described the task of automating delivery of guided expert help. To reach thousands of people using trained professionals is a daunting and expensive proposition. Our goal was to serve tens of thousands of people with little direct human intervention. We reasoned that by employing big data analytics, some aspects of artificial intelligence, and machine learning (selectionist) algorithms, we could design a program that interacted with a user much the same as would a human counselor. And perhaps eventually—as the program learned more and more—even better. In Part 1 I described the technologies available. In Part 2 I described how those technologies could be used to assess and prescribe strategies and goals tailored to the individual user. This final entry discusses how we approached designing an app that would guide the user through their program and enable them to reach their critical goals.

The Problem

Many of us have used or at least seen clever apps available to assist us in meeting behavioral goals. These apps gather data—from user input, sensors, or GPS—and present that data in ways designed to help us meet our goals. Some help set goals, some provide advice, some connect us to support communities, and some allow us to consult with experts or more knowledgable people. Some apps are designed to be used in therapy, allowing clients to record their thoughts, feelings, and actions and later send them to a therapist. Although data are collected and analyzed and some advice or activities may be supported through audio or video coaching, most interventions are left to human coaches or therapists. While all of this may be helpful, the cost of providing a personalized, therapeutic intervention at scale remains high. Would it be possible to automate clinical intervention?

The Solution

The solution we designed would involve using commercially available machine learning software whose adaptive algorithms would allow us to pinpoint the complex interrelationship of multiple variables that influence behavior and produce highly accurate recommendations and results. We would use these algorithms and other artificial intelligence-based software to transform big data and our Constructional Questionnaire into personalized goals and subgoals that we hoped would increase enrollment and engagement of large populations. We then reasoned we could use the same underlying adaptive algorithm to create personalized plans that incorporated health behaviors into the consequentially important patterns of an individual's daily life.

The app would have to generate individual target repertoires or goals for users based upon a combination of big data analysis and the Constructional Questionnaire. These goals would be broken into subgoals based on each person's current relevant repertoire. Each week's subgoals would be based on the past week's performance and the analysis of the contingencies responsible for that performance. We wanted to begin with the leanest intervention possible that would allow our users to meet their subgoals, which would ultimately lead to meeting their target goals.

Our model was Goldiamond's (1974, 1975) constructional approach to self-control. In this approach, weekly subgoals are typically determined after a constructional therapist and the client analyze daily logs kept by the client. Each subgoal is chosen based upon its relation to the final goals. The daily log is often the key to success. Using the logs, clients can identify if they are getting what they want out of their daily interactions and activities. They can determine if the reinforcers important to them are likely to occur as a result of what they are doing. They can test different approaches and adjust their behavior in order to achieve what they value. Clients make an entry every hour on average. An example of such a log is provided below.

No.	Time & Duration	Audience, Place, Conditions	Activity Intended, What I Wanted	Activity, What I Got	Comments, Emotions

Of great significance is the distinction between what was wanted versus what actually transpired (“What I Got”). “What I Wanted” speaks to the potentiating variables operative at that time. That is, what the likely reinforcers were at that point and why they are important. These may be either positive reinforcers (“Him to ask me out”) or negative reinforcers (“Mom to stop nagging me about homework”). Emotions are used as windows into the contingencies operating. Through a series of questions and a dialogue with the client, instances are cumulatively examined and analyzed until next steps emerge. For example, we may discover that the only real interaction with mom comes in the form of nagging and no instances of close interaction were recorded. This suggests that the consequences of nagging and what occasions nagging may be critical to understanding the contingencies of which the overall pattern is a function. That is, when combined with what is in the log, what isn’t in the log may be as important as what is in it. We may find that although the nagging is aversive to a certain extent, it is also reinforcing, bringing close social contact not otherwise available. The question would be raised, “How can a valued and meaningful interaction with mom be developed without the behaviors that occasion nagging?” The logs would be examined to look for a place to start—where a brief, meaningful interaction could occur. The interaction would be a subgoal for the next week and the results evaluated during the next session. In the next week’s logs we might see:

Activity Intended, What I Wanted: A nice interaction with mom.

Activity, What I Got: Asked mom about how she picks out handbags since she always looks so great. She explained it was a balance between utility, activity, and matching ones’ clothes.

Comments, Emotions: Felt really good to have these minutes just talking, I almost cried.

And in a later entry: I noticed she nagged me less, too. Maybe this is how we have learned to interact with each other.

But how can an app facilitate this type of analysis and interaction? From the constructional assessment and big data analysis, possible activities and “wants” related to possible circumstances (audience, place, conditions) could be determined through the application of machine learning techniques not unlike those described in Part 2 of this series. The circumstances could be built from a menu of items proposed by the app. Another field would allow manual entry, which could later be incorporated into the menu structure. Once these were determined, possible activities intended would be proposed from which the client would select. Again, these would be generated by not only current inputs, but also past inputs and other big data sources, both individual and group, and by evaluating how close the proposed scenario was to the actual activity intended. Further, the Comments, Emotions field could be used to suggest the consequential contingencies operating. For example, after the entry on mom nagging, an emotional response might be “irritated, but want closeness.” “Want closeness” could be identified as indicating that there may be a positive reinforcer operating and produce a clarifying question that would not otherwise occur.

If the program worked properly, it would be a therapist that would continually get better and better. No confirmation bias or other response bias would enter into the therapeutic intervention.

As the software is used over time, variations in activities based on the hourly entries and their commonalities and differences would evolve until the processes of offering log entries for each cell became more accurate and linked to what the client is after. These data would use the emotional comments (selections) as indicators of the consequential contingencies operating and machine learning algorithms would increasingly relate those to the log entries and be used to suggest subgoals, which would themselves go through a rating system, the results of which would be fed back to the selectionist algorithms operating throughout the program.

Mini scenario evaluations would be presented and evaluated by the client. This would include contingency observations and suggestions. These in turn would be rated or ranked. In essence, the selectionist algorithms and the proposed suggestions would take on the role of a human therapist. If the program worked properly, it would be a therapist that would continually get better and better. No confirmation bias or other response bias would enter into the therapeutic intervention.

The results from the continual evaluation would also inform the scenarios of the original Constructional Questionnaire and be included in the big data analysis both individually and for all users in order to evolve better assessments and planning. The goal was to build a system that continually improves itself the more it is used and the greater the number of people who used it.

…behavior analytic principles could potentially be the most important contributor to a successful automated clinical intervention system.

In summary, the app would not only provide highly valued client goals and the initial and subsequent subgoals based upon an analysis of the logs, it would suggest steps or strategies to get there that would have their initial origin in the constructional interview, big data analysis, and log entries. The outcome of each highly ranked (by the user) step or strategy would enter the selectionist (machine learning) big data database and be used to further refine the program.

To accomplish this, we would use the machine learning software to produce highly accurate recommendations and results. We would use big data to help determine actionable insights to increase enrollment and engagement of large populations. We would then create personalized plans that incorporated health behaviors into the consequentially important patterns of an individual's daily life.

If, after a time, progress stalled, or app usage began to drop off, that would occasion a reassessment. Something important to the client may have been missed or benefits of the disturbing pattern not entirely identified. Given the patterns we were targeting and the types of possible interventions required as described in Part 1, we assumed many interventions would require a more topical approach. Often, however, a systemic rather than topical intervention might be required.

Quality behavioral interventions can conceivably be delivered to millions, with very few actual therapists involved.

In a systemic intervention, a matrix of contingencies is identified that typically does not contain the presenting complaint, yet when addressed the disturbing pattern may drop out. An example would be a person seeking help with weight reduction who leaves their desk at work to eat a candy bar when "stress" and "anxiety" build. It may be determined that leaving the desk and eating the candy provides a needed break from the demands of coworkers and gives the client time to gather his thoughts. Stress implies increasing work requirements with falling reinforcement rates, and anxiety suggests behavioral requirements for which the client may be unprepared to meet. The program would scan the daily log for indications that this is happening. In particular, the daily logs would be examined to determine how the client controls work requirements and how the client evaluates work requirements and prioritizes them. The app might then recommend a program of assertiveness training and another focusing on organizational skills as part of the subgoals. No action in regard to the candy eating would be taken. Once the new assertive and organizational patterns were established and greater control over the work environment occurred, the candy eating should no longer be required and drop out. The app would continually scan the logs for evidence that the frequency of candy eating changed. Such a reduction in candy eating frequency would suggest the systemic procedures worked. As the program made recommendations and adjusted to user success (or lack thereof), its recommendations should get better and better. Further, individual user data would be continually aggregated with others to look for procedures that, given certain circumstances, would be the most successful. These data would also be fed back to the big data analysis and assessment algorithms in order to improve them.

We were confident we could, within one or two years, build such a system. And, behavior analytic principles could potentially be the most important contributor to a successful automated clinical intervention system. As I said at the outset of this series, we didn't do it. But, I firmly believe it could be done. The technology and tools exist. The implications for treatment and service delivery are immense. Quality behavioral interventions can conceivably be delivered to millions, with very few actual therapists involved. And further, as the user database grows and the data are fully integrated and related, the automated therapist may grow to be much more effective than a human therapist. I believe that sooner than many realize, the future of clinical behavior analysis likely lies with the machines.

References

Goldiamond, I. (1974). Toward a constructional approach to social problems: Ethical and constitutional issues raised by applied behavior analysis. *Behaviorism, 2,* 1–84. (reprinted 2002 in Behavior and Social Issues, 11, 108-197.

Goldiamond, I. (1975). A constructional approach to self control. In A. Schwartz & I. Goldiamond (Eds.), *Social casework: A behavioral approach* (pp. 67–130). New York: Columbia University Press.

Goldiamond, I., & Dyrud, J. E. (1968). Some applications and implications of behavioral analysis for psychotherapy. In J. M. Shlien (Ed.), *Research in psychotherapy* (Vol. 3, pp. 54–89). Washington, DC: American Psychological Association.

Layng, T. V. J., Andronis, P. T., Codd III, R. T., Abdel-Jelil, A. (2022). *Nonlinear contingency analysis: Going beyond cognition and behavior in clinical practice*. NY: Routledge.