Intelligent Agents with Subjective Experience

Ana Pasztor (PASZTORA@CS.FIU.EDU)
School of Computer Science
Florida International University
University Park, Miami, Fl 33199

Abstract

In this paper I am presenting a model of human communication which allows a decomposition of our subjective experience into units called strategies. Strategies are systematically ordered sequences of representations in our sensory modalities, endowed with so called submodality distinctions, orientation, effect, and links to other representations. As certain people prove more successful in performing given tasks than others, we want to model their strategies and transfer them to other intelligent agents. Unlike most philosophers who have written about qualia and believe that subjective experience condemns us to "solitary confinement," I believe strategies are accessible to us. Here is how. All behavior, from learning to communication and change, is the result of our strategies. It is constantly giving us information about the underlying strategies by way of physiological and linguistic cues. The model I am presenting in this paper offers a systematic study and categorization of these cues. Utilizing these cues, we can now elicit people's successful learning strategies, and design psychologically plausible intelligent agents, which learn effectively and behave consistently with the subjective experience we want them to have, or, even make them have.

The Units of our Subjective Experience

I share with Dilts *et al.* (1980) the belief that "any portion of our experience can be understood and eventually controlled if we are willing to study the processes which underlie that experience."

Such control can be achieved by decomposing the processes underlying our subjective experiences into *units* that can be readily changed, learned, and transferred from one intelligent agent (including humans) to another, in order to help achieve desired goals.

Another way to put it is using a metaphor borrowed from (Cameron-Bandler & Gordon & Lebeau, 1985): we humans are most sophisticated computers, which, alas, come to this world without manuals. So it is entirely up to us to figure out how we work, particularly which software to use or not to use for a given task. Some people have the good fortune to be using programs which are optimal for certain tasks (Boden, 1990; Weisberg, 1988; Dilts, 1994). In this paper I will present a model of human communication, called Neuro-Linguistic Programming (NLP henceforth), the goal of which is to study these people's programs and chunk them down into transferrable units, so that we will be able to *model* them and eventually have any intelligent agent utilize their programs.

NLP evolved from the work of family therapist Virginia Satir, Gestalt therapy founder Fritz Perls, hypnotherapist Milton Erickson, anthropologist Gregory Bateson, and many others. It provides a model of understanding subjective experiences, as well as finding and using cues to interpret these subjective experiences (Dilts & Dilts & Epstein, 1991).

All information processing happens through our senses. All subjective experiences are manifested (consciously or unconsciously) in our sensory (representational) systems: visually, auditorily, kinesthetically (for the sake of simplicity, I include here emotions, external and internal feelings, as well as taste and smell). Let us denote them by **V**, **A** and **K**, respectively.

Talking about unity of consciousness, Searle (1992) says:

> Conscious [and I argue that all mental] states . . . come to us as part of a unified sequence. I do not just have an experience of a toothache and also a visual experience of the couch that is situated a few feet from me and of roses that are sticking out from the vase on my right, in the way that I happen to have on a striped shirt at that same time as I have on dark blue socks. The crucial difference is this: I have my experiences of the rose, the couch, and the toothache all as experiences that are part of one and the same conscious event. Unity exists in at least two dimensions, which, continuing the spatial metaphors, I will call 'horizontal,' and 'vertical.' Horizontal unity is the organization of conscious experiences through short stretches of time.... Vertical unity is a matter of the simultaneous awareness of all the diverse features of any conscious state, as illustrated by my example of the couch, the toothache, and the rose. We have little understanding of how the brain achieves this unity. . . . Without these two features - the horizontal unity of the remembered present [Edelmann] and the vertical unity of the binding of the elements into a unified column - we could not make normal sense of our experiences...

Vertical Units of our Subjective Experience

In sticking to the (NLP) notation used above, the elements of our subjective experience which are responsible for the vertical unity that Searle was talking about - we will call them *vertical units* of our experience - can be denoted -see (Markova, 1992) by, say,

$$\begin{pmatrix} V \\ A \\ K \end{pmatrix}:$$

At each moment in time our subjective experience is manifested/represented through the three modalities V,A, and K, where the top one denotes the representational mode we attend to most consciously, like visual in the above example - e.g., when we concentrate on the shape of the rose, and the bottom one is that which we attend to least consciously at that moment, like kinesthetic in this case.

Sometimes we may have a synesthesia of two or more representations (e.g., see the smell), in which case we can fuse together two representations (e.g., V/K).

In mainstream NLP only the most dominant modality of the vertical unit is taken into consideration, a practice which I will also follow below by referring to the units simply as *representations*.

The great importance of the kinesthetic component of each state/vertical unit, especially of the physical aspects of this component is best explained in Damasio's (1994) "somatic-marker hypothesis." He creates an imaginary scenario, which calls for sudden decision making. "The brain of a normal, intelligent, and educated adult reacts to the situation by rapidly creating scenarios of possible response options *and* related outcomes. To our consciousness, the scenarios are made of multiple imaginary scenes, not really a smooth film, but rather pictorial flashes of key images in those scenes, jump cut from one frame to another, in quick juxtapositions." (Please note the qualities of the images — c.f. below, where I discuss submodalities.) The mind excites "your consciousness in a show too rich for you to encompass fully. Even in this caricature you will recognize the sort of quandary we face most every day. How do you resolve this impasse? How do you sort out the questions inherent in the images before your mind's eye?" You follow your "gut feelings," or your "feeling or rightness," which mark your images and are felt in your body, hence the name "somatic-marker."

Each representational system is designed to 'perceive' certain basic qualities of the experience it represents. These include characteristics such as color, brightness, tone, loudness, temperature, pressure, etc. These qualities are called **sub-modalities** in NLP, since they are sub-components of each of the representational systems.

Sub-modalities play an important role in determining the meaning of our experiences. For example, think of a pleasant experience you had. Now slowly change sub-modalities of your memory like (visual:) color, brightness, focus, size, movement etc.; (auditory:) pitch, melody, rythm etc.; (kinesthetic:) movement, duration, speed, tempo etc. Notice whether and how each change in sub-modality changes the meaning of your memory. Another example: when asked to describe the cognitive processes involved in e.g. *anticipating* versus *hoping* for an event, most people report that in anticipating they have a steady image of the event, while in hoping they experience a mental flipping back and forth between an image of the event they were hoping to happen and an image of the event they hoped will *not* happen. By holding still the image of the event hoped for, they experienced anticipation instead of hope. Similarly, it is impossible to maintain an emotion of impatience and at the same time experience a slow *tempo*.

In view of these submodalities, somatic-markers are those *components* of our kinesthetic representations, which, I argue, are our real decision making agents.

A sensory representation also has **orientation**, which can be *external*, like in the case of perceptions, or *internal*. Internal representations can be *remembered* or *constructed*. Auditory representations can also be *digital*, like in the case of words.

> Another general feature of each modality is that it can occur under the aspect of PLEASANT or UNPLEASANT, and the way in which it is pleasant/unpleasant is in general specific to the modality. ...Often but not always, the pleasure/unpleasure aspect of conscious modalities is associated with a form of intentionality (Searle, 1992).

Positive (pleasant) kinesthetic representations are denoted by K^+, while negative (unpleasant) ones by K^-. Each representation or vertical unit has an **effect**. For example, an effect could be the function to generate or input a sensory representation, or to change or test/evaluate some part of our experience.

Horizontal Units of our Subjective Experience

Coming back to Searle's *horizontal unity* of consciousness, this is achieved by sequences of representations or vertical units which have consolidated into behavioral units, called strategies (Dilts *et al.*, 1980). Analyzing a letter Mozart wrote in 1789 describing his strategy for composing music, Dilts (1994) outlines the first, most basic stage in this strategy as follows:

> Mozart's composition process seems to begin from a positive, congruent kinesthetic feeling state combined with some kind of physical movement [K^+/K^e]. Internal feelings [K^i] arise from this state which produce sounds or tones through a natural overlap between the kinsthetic and auditory senses ($K \rightarrow A^i$). This linking is known as a *"syntesthesia."* If these sounds fit with or reinforce the positive feeling state [K^+] they are hummed (K^e/A^e) and retained in memory, otherwise they are discarded.

As we can see in this example, each representation/vertical unit of our experience serves as **anchor** for the next representation/vertical unit. For example, I hear X's name (A^e_d), this triggers my remembering X's face (V^r), close up, somewhat distorted, and pinkish red, which, in turn, triggers a negative emotion (K^-). Now I also could hear someone say X's name (A^e_d) in a really small, squeaky Micky-Mouse type of voice, which would perhaps trigger no image of X's face at all, or perhaps trigger a small image of X against some background (V^i), which, in turn might trigger a funny feeling (K^i).

According to (Dilts *et al.*, 1980), "A basic assumption behind anchoring is that all experiences are represented as gestalts of sensory information ... Whenever any portion of a particular experience or 4-tuple [vertical unit] is reintroduced, other portions of that experience will be reproduced to some degree. Any portion of a particular experience then may be used as an anchor to access another portion of the experience."

The link from a representation (or vertical unit) of our experience at a time point *t* to a representation (or vertical unit) of our experience at a time point *t+1* depends hereby on anchoring, which in turn depends on subjective, social, cultural and many other factors, perhaps also captured in Searle's (1992) notion of Background.

Making Agents Believable

Now we have a model of the structure of subjective experience. However, in order to transfer subjective experience from one agent to another, or to build believable agents (Bates, 1994), we need to be able to *recognize* this structure in ourselves and in others.

Sequences of images go by so very fast, that people usually are not aware of them. If you ask someone who had pizza for lunch "how do you know that you didn'ty have quiche for lunch?" they will often say "because I remember," but they will rarely say something like "the qualities (submodalities) of the images involving the pizza are different from those involving the quiche, and according to my reality strategy it is obvious I had pizza." But introspection with awareness of the representations which make up our subjective experiences can be learned and taught, and reveals a very rich structure of these experiences, which we used to think of as being atomic - c.f. (Pasztor, 1996).

Denett (1988) seems to agree:

> This putative grainlessness, I hypothesize, is nothing but a sort of functional invariability: it is close kin to what Pylyshyn (1980,1984) calls *cognitive impenetrability...* so what counts for an individual as the simple or atomic properties of experienced items is subject to variation with training . . . Consider the results of 'educating' the palate of a wine-taster, or 'ear training' for musicians. What had been 'atomic' or 'unanalysable' becomes noticeably compound and describable....

Much more widely disputed is our ability to understand *other* people's ("third person's") subjective experiences.

Here is what Dennett (1988) has to say:

> ... far from being directly or immediately apprehensible properties of our experience, they are properties whose changes or constancies are either entirely beyond our ken, or inferrable (at best) from 'third-person' examinations of our behavioral and physiological reaction patterns....

The elements of the NLP model I presented in the previous section "deal primarily with cognitive processes. In order to function, however, these mental programs need the help of certain bodily and physiological processes for consolidation and expression. These physical reactions are important for the teaching or development of certain mental processes as well as for the external observation and confirmation of them. The primary behavioral elements involved . . . are: **Body Posture**, **Accessing Cues**, **Gestures**, **Eye Movements**, and **Language Patterns**" (Dilts & Dilts & Epstein, 1991).

Below are some examples from (Dilts & Dilts & Epstein, 1991).

Leaning back with head and shoulders up or rounded, shallow breathing indicates the use of visual, *body leaning forward, head cocked, shoulders back, arms folded* indicates the use of auditory, while *head and shoulders down, deep breathing* indicates the use of kinesthetic representations.

According to Dilts and Dilts and Epstein (1991), "When people are thinking, they cue or trigger certain types of representations in a number of different ways including: breathing rate, non-verbal 'grunts and groans,', facial expressions, snapping their fingers, scratching their heads, and so on. Some of these are idiosyncratic to the individual and need to be 'calibrated' to the particular person. Many of these cues, however, are associated to particular sensory processes."

High shallow breathing, squinting eyes, voice higher pitch and faster tempo are some visual, *diaphragmatic breathing, knitted brow, fluctuating voice tone and tempo* are some auditory, while *deep abdominal breathing, deep breathy voice in a slower tempo* are some kinesthetic cues or triggers.

According to the NLP model, eye movement is not at all random. There are patterns in our eye movement, which, like other physiological patterns mentioned above, reveal the structure of our subjective experience. For details see (Dilts & Dilts & Epstein, 1991). NLP's claim about eye movement has been subject to numerous experiments in the 80s, but no conclusive results have been achieved (see next section). Since this is a topic of utter importance to areas of AI such as building believable intelligent agents (Bates, 1994), or Participatory Design (Beyer & Holtzblatt, 1995), I would like to urge the research community to continue the experiements of the 80's with the modifications suggested by Einspruch & Forman (1985) - see next section.

Regarding gestures, "People will often touch, point or use gestures indicating the sense organ they are using to think with" (Dilts & Dilts & Epstein, 1991). Thus *touching or pointing to the eyes and gestures made above eye level* indicate the use of the visual, *pointing toward or gesturing near the ears and touching the mouth or jaw* indicate the use of the auditory, and *touching the chest and stomach area and gestures made below the neck* indicate the use of the kinesthetic modality.

Most of the trouble in human communication, I believe, comes from our failure to use the enormous richness of both non-verbal and verbal cues to our subjective experience in a conscious and systematic way. As far as recent research on these cues goes, there is far too little out there. There is some on gestures, for example, for the purpose of computer animation - see (Cassell *et al.*, 1994; Bates, 1994; and references therein).

Besides cues, one can use questions to elicit a person's subjective experiences. Taking introspection and self-report seriously, Bandler and MacDonald (1988) have implemented a list of *submodality distinctions* in each representation system, accompanied by possible questions that can be used to elicit the respective distinction. Here is a sample. (Each submodality distinction is going to be followed in parentheses by some questions to elicit that distinction.)

(**Visual**:) Color/black and white (Is it in color or black and white? Is it full color spectrum? Are the colors vivid or washed out?); Brightness (In that context, is it brighter or darker than normal?); Focus (Is the image sharp in focus or is it fuzzy?); Detail (Are there foreground and background details? Do you see the details as part of a whole or do you have to shift focus to see them?); Size (How big is the picture? - ask for specific, estimated size, like 11" x 14"); *Movement* within the image (Is it a movie or a still picture? How rapid is the movement: faster or slower than normal?); Associated/dissociated (Do you see yourself or do you see the events as if you were there?) etc.

(**Auditory**:) Location (Do you hear it from the inside or from the outside? Where does the sound (voice) originate?); Pitch (Is it high-pitched or low-pitched? Is the pitch higher or lower than normal?); Volume (How loud is it?); Tempo (Is it fast or slow?); Rhythm (Does it have a beat or a cadence?); Duration (Is it continuous or intermittent?) etc.

(**Kinesthetic**:) Quality (How would you describe the body sensation: tingling, warm, cold, relaxed, tense, knotted, diffused?); Intensity (How strong is the sensation?); Location (Where do you feel it in your body?); Movement (Is there movement in the sensation? Is the movement continuous or does it come in waves?); Direction (Where does the sensation start? How does it get from the place of origin to the place where you are most aware of it?); Speed (Is it a slow steady progression or does it move in a rush?) etc.

In the NLP model presented in this paper, metaphors used by people to express their subjective experiences are understood as cues to the strategies and submodality distinctions they are using. (Very few researchers take linguistic metaphors seriously - see (Lakoff, 1987)).

Below is a sample of such verbal cues (c.f. Bandler & MacDonald, 1988).

(**Visual**:) Things were blown out of proportion; My job seems overwhelming; Life is so drab; I need some distance from it; He's had a colorful past; That throws a little more light on it; It all seems so hazy; I don't know, it just flashed on me; When you said that I just saw red; That brightens up my day; etc.

(**Auditory**:) The right decision was screaming at me; She gives me too much static; It's just a whisper; If I nag myself long enough, I'll do it; There is too much discord in our relationship; We need to orchestrate our vacation; I keep telling myself, "You can't do anything right;" etc.

(**Kinesthetic**:) It's got a slimy feel to it; He's hot; She's a cold fish; Whenever I hear that, my stomach knots up; The pressure is off; The whole thing weighed on my mind; I'm off center, like everything is out of kilter; I'm trying to balance one against the other; Yeah, I feel up to it; etc.

Background Research

NLP is a model of human communication, that fares excellently when tested by general criteria used for evaluating cognitive theories, such as representational and computational power, psychological and neurological plausibility, and practical applicability - see (Thagard, 1996).

But there has been virtually no connection made to cognitive science. A few years ago Christine Lisetti and I have completed a study of 49 papers written on the NLP model. Most papers were written in the 80s. The findings of 16 of these papers were negative or questioning, while the findings of the others were positive. However, the negative or questioning findings were mostly based on experiments which were conducted improperly. The comments of Einspruch and Forman (1985) give a very acurate analysis of these experiments, as well as suggestions on how proper experiments could be conducted. To better understand the comments, it is necessary to mention, that in the 80's NLP was still almost exclusively being applied in therapy.

> There is a growing body of empirical literature on NLP. A review of this literature by Sharpley1984 failed to consider a number of methodological errors. In the present article the authors identify six categories of design and methodological errors contained in the 39 empirical studies of NLP documented through April 1984. ...It is very difficult to understand the NLP framework from the perspective of traditional counseling models; it is much *more appropriate to approach it from the framework of mathematics, biology, or cybernetics.* ...NLP is a complex model requiring extensive training before a practitioner may legitimately undertake a study of this nature. A number of **modifications** could be made to improve designs of research conducted on NLP.

They then go on to list these modifications.

In discussing Flanagan's contribution to the 25th Carnegie Symposium on Cognition, Galin (1993) says: "He [Flanagan] concluded that a scientific and unified theory of consciousness is possible, but that what it will look like depends on who you let contribute to it. For a complete theory he suggests that the cognitive functional analysts, the neurobiologists, and those interested in phenomenology (subjective experience) should all be invited. In addition to cognitive psychology and neuroscience, he says, 'give phonomenology its due. Listen carefully to what individuals have to say about how things seem.' But at this meeting there has been enormous resistance to talking about subjective experience, even though we are supposed to have outgrown the Behaviorists' restriction to external observation and their edict forbidding introspection. ...Well, first-person experience is politically correct now; what is the reason that we still are not talking about it?"

Galin (1993) suggests that one reason for this is a lack in "fundamental cognitive tools for the task", a lack in taxonomy and in "consensus on the constructs or terminology with which to characterize its richness ...."

Although NLP provides more than just the fundamental cognitive tools for studying subjective experiences, it still needs to be incorporated into scientific research. I agree completely with Galin (1993) when he says "Research on subjective experience is like sex; everybody is terribly interested in it but does not want to be caught in public doing it. Could it be that this area of research contains something so personal, so intimate, that we have trouble looking at it directly? Let's peek."

Applications

Number 7 of *Communications of ACM*, vol. 37, is dedicated to Intelligent Agents. All authors agree that there is a need for agents with common sense, agents capable of learning (see e.g., Mitchell *et al.*, 1993; Maes, 1993). In the abovementioned volume, in *A Conversation with Marvin Minsky About Agents*, Minsky notes:

> The problem with this, of course, is that we still don't have suitable learning techniques. Lenat and I agree in the view that in order to learn as a person does, one will need to begin with a considerable body of built-in knowledge about a variety of effective ways to learn. The problem is we have not done enough research yet to know how to do this.

Our strategies, i.e. the representational modalities we use to represent our experiences, their orientation, submodalities, effects, and the links between them (anchors), are the software programs we use to go through our lives. Certain programs are more suitable for certain tasks than others. In order to acquire the "knowledge about a variety of effective ways to learn," we need to study various learning strategies, stuccessful and unsuccessful. The model I presented above, allows us to do just that. Here is an example.

Diane Lutzinger (1991), a teacher trainer, gives the following account of a student's unsuccessful vs. successful strategies.

> When Jason came to Diane Siegel for help, the first step was to find out the specific behaviors he used in class. For example, Jason would listen to the teacher explain a history lesson, then he'd have a little conversation in his head that went like this: 'This is stupid. Why do I need to know the date of this battle or the name of that treaty? I'll never remember it anyway. The teacher is so boring!' Without being aware of it, Jason would then feel bad about the lesson, the teacher and himself. Feeling bad, he would shut down his attention to what was going on in class, and the cycle would continue its downward spiral.

In contrast, here is Jason's football strategy:

> The coach would describe the play. Jason listened and imagined the formations and tactics as he remembered them from 'chalk talk' when the coach drew figures on the board. His internal conversations went more like this, 'I know this play. I can see which player will run out for the pass and which blockers will move forward to assist. I'm able to do this play. I'm excited. Let's get 'em!

Jason's history strategy is "listen, self talk, feel bad, loop back to listen," while his football strategy is "listen, picture, self-talk, feel good" (Lutzinger, 1991).

Once strategies, be they learning or otherwise, become cognitively penetrable and disclosed, they become "programmable, hence subject to change" (Gardner, 1985). Both representational modes, with their orientations, submodalities, and effects, as well as their links, can be changed, as extensively explained and demonstrated by Dilts *et al.* (1980).

In setting up evaluation criteria for cognitive theories/models, Thagard (1996) lists three milestones of **applicability**: *education, design, and intelligent systems.*

I have dealt with applicability in *education* above. As far as design goes, the structure of subjective experience and its cues play a very important role in Participatory Design (Namioka & Schuler, 1993), where an excellent communication between all participants in the design process is highly critical, as well as in modeling creativity (Boden, 1990; Weisberg, 1988; Dilts, 1994).

As an example, I will mention a creativity strategy developed by Dilts (1994). Dilts modeled it after Walt Disney's strategy to create his animated movies with his team. The strategy is best described as a loop which goes from a dreamer phase to an implementation phase to a critique phase and back to a dreamer phase etc., until the critic in the critique phase is satisfied. The phases are clearly separated so they do not "contaminate" each other (have you ever experienced the voice of an internal critic while you were trying to come up with an idea?), moreover each phase builds upon and improves the previous one.

As far as *intelligent systems* go, particularly Intelligent Agents, the applications are numerous.

Intelligent Agents have to communicate effectively with their users, as well as with each other. *Rapport* with a partner is the key to effective communication, and can be achieved by closely observing and responding adequately to *verbal and non-verbal cues* to how the partner experiences the world. Maes (1993) talks about *trust* (i.e. how we can "guarantee the user feels comfortable delegating tasks to an agent") as a main problem "to be solved when building software agents." Mitchel *et al.* (1993) agree that "user acceptance will depend on customizing the personal assistant to the particular habits and interests of the user." Trust and user acceptance, on the other hand, begin with believability of the agents. "Emotion is one of the primary means to achieve this believability, this illusion of life, because it helps us know that characters really care about what happens in the world, thay they truly have desires" (Bates, 1993).

Cues to our experience are the "regularities and recurrent patterns" Maes (1993) talks about when she says "the interface agent learns by continuously 'looking over the shoulder' of the user as the user is performing actions," the goal being to *find and automate* these cues.

I have presented in this paper a model of human communication which allows us to recognize the building blocks of subjective experience by way of body posture, accessing cues, gestures, eye movements, and language, and utilize them to build intelligent agents we trust and accept. Moreover, I suggest that we can even go a step further, and not only build intelligent agents which *seem* to have subjective experiences, but also build such which actually have them.

References

Bandler, Richard & MacDonald, Will (1988). *An Insider's Guide to Sub-Modalities*. Meta Publications.

Bates, J. (1994). Agents and Creativity, *Communications of the ACM*, 37(7), 122–125.

Beyer, H.R. & Holtzblatt, K. (1995). Apprenticing with the Customer. *Communications of ACM*, 38(5), 45-52.

Boden, M.A. (1990). *The Creative Mind: Myths and Mechanisms*. Basic Books.

Cameron–Bandler, Leslie & Gordon, David & Lebeau, Michael (1985). *The Emprint Method. A Guide to Reproducing Competence*. Moab, Utah: Real People Press.

Cassell, Justine & Stone, Matthew & Douville, Brett & Prevost, Scott & Achorn, Brett & Steedman, Mark & Badler, Norm & Pelachaud, Catherine (1994). Modeling the Interaction between Speech and Gesture. In Ashwin Ram & Kurt Eiselt (Eds.), *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society, Atlanta, Georgia, August 1994*, (pg. 153).

Damasio, Antonio R. (1994). *Descartes' Error*. Grosset/Putnam.

Dennett, Daniel C. (1988). Quining qualia. In Marcel, A.J. & Bisiach, E. (Eds.), *Consciousness in Contemporary Science*, pp. 42-77. Oxford: Clarendon Press.

Dilts, R.B. (1994). *Strategies of Genius*. Meta Publications.

Dilts, R.B. & Dilts, R.W. & Epstein, T. (1994) *Tools for Dreamers; Strategies for Creativity and the Structure of Innovation*. Meta Publications.

Dilts, R. & Grinder, J. & Bandler, R. & Bandler, L. & Delosier, J. (1980). *Neuro-Linguistic Programming: Volume I, The study of the Structure of Subjective Experience.*, Cupertino, Calif.: Meta Publications.

Einspruch, E. L. & Forman, B. D. (1985). Observations Concerning Research Literature on Neuro–Linguistic Programming. *Journal of Counseling Psychology*, 32(4), 589–596.

Galin, David (1993). What is the Difference between a Duck?, In Cohen, Jonathan & Schooler, Jonathan (Eds.), *Scientific Approaches to Consciousness: 25th Carnegie Symposioum on Cognition, May 20-22, 1993*. Hillsdale, NJ: Erlbaum.

Gardner, Howard (1985). *The Mind's New Science*. Basic Books.

Lakoff, George (1987). *Women, Fire, and Dangerous Things*. The University of Chicago Press.

Famiano Lutzinger, Diane (1991). Smart Stuff. 'Maps' of Learning. *Anchor Point* June '91, 14- 16.

Maes, P. (1994) Agents and Creativity. *Communications of the ACM*, 37(7), 31–146

Markova, Dawna (1992). *How your child is smart: a life changing approach to learning*. San Bernadino, CA : R. Reginald ; Borgo Press.

Mitchell, T. & Caruana, R. & Freitag, D. & McDermott, J. & Zabowski, D. (1994). Experience with a Learning Personal Assistant. *Communications of the ACM*, 37(7), 81-91.

Namioka, A. & Schuler, D. (1993). *Participatory Design: Principles and Practice*. Hillsday, NJ: Erlbaum.

Pasztor, A. (1996). Qualia Decomposed. *Consciousness Research Abstracts: Toward a Science of Consciousness 1996 (Special volume of Journal of Consciousness Studies* (pg. 41). Imprint Academic.

Polya, George (1945). *How to Solve it*. Princeton University Press.

Rieken, D. (1994). A Conversation with Marvin Minsky about Agents, *Communications of the ACM*, 37(7), 23–29.

Searle, John, R. (1992). *The Rediscovery of the Mind*. MIT Press.

Sharpley, C. (1984). Predicate Matching in NLP: A Review of Research on the Preferred Representational System. *Journal of Counseling Psychology*, 2(31), 238–248.

Thagard, Paul. (1996). *Mind. Introduction to Cognitive Science*. The MIT Press.

Weisberg, R.W. (1988). Problem Solving and Creativity. In Sternberg, R.J. (Ed.), *The Nature of Creativity*. Cambridge University Press.