1 An Essay on Part of the AHI Program

1A J. F. Rulifson, 8 June 1968

1B .NBL=2;

1C Foreword

1C1 This paper is a term project for Phil. 245 (Psycholinguistics), Spring Qtr. 1968. The Foreword discusses the origin and history of the thoughts presented later. The paper then loosely describes the practical, architectural goals of the NLS (On-Line System). The currently implemented NLS is subsequently analyzed from a more general, formal viewpoint. The notion of "intellectual augmentation" is discussed. A general strategy is offered for testing, evaluating and designing system features. Finally, two possible experiments are outlined.

1C2 The impetus to write a paper such as this grew out of three strong, personal feelings.

1C2A The AHI Center is supposed to "Augment Human Intelligence" but no one can give a precise explanation of what that means. There is, therefore, no measure of success and this results in a personal feeling of through frustation.

1C2B Research in the AHI Center is at a standstill, and the rut is impossible to get out of unless a fresh approach is

taken to NLS.

1C2C After all the work that has gone into building a flexible system, I want to see the flexibility used to truly improve the system. There is a wealth of potential. It can be exploited only if we understand the potential, the process being augmented, and the interconnection.

1C3 The relationship of this to psycholinguistics will be seen later. NLS is supposed to be a system that augments its user's capacity to organize and remember information. Language plays two roles, so important that they led me to Phil. 245 in search of ideas.

1C3A The user must talk to the computer. No matter what he does, he is using some form of language, even though the language may be trivial.

1C3B Part of the information which the user manipulates and organizes will be stored in natural language.
Psycholinguistics bears heavily on determining what the user does or accomplishes through the transcription and manipulation of the information.

1C4 The paper grew in the following way:

1C4A During November 1967 I happened on some of Piaget's work

concerning groupements. I became very excited and tried to relate his work to NLS. A term paper for CS225 summed up these results--it was a complete disaster. The vagueness of the then NLS framework coupled with my rudimentary understanding of Piaget resulted in groping thoughts and nothing concrete.

1C4B To try and get some direction I talked with Colby, and discussed the possibility to doing some kind of testing with the current system. I then felt that only through some sort of empirical study could new, fresh thought come forth. Colby suggested that I look at an experiment designed to discover the way small children index limited, self-contained fields of information. The idea was appealing but it was not tied to the formework of the project. Slwoly I developed the attitude that there was no adaquate framework to couch an experiment in.

1C4C In early May Suppes convinced me that empirical testing must be done in parallel with clinical work. This made some problems very clear, for statistical studies require a concrete task, or problem domain, and I could not conceive of one that was related to the overall goals of NLS. What was needed was an explicit formalism that NLS could be put into. Statistics could then test statements about the formalism.

1C4D This paper now began to take form. The notions of separating the control language from the forma (defined below)

were laid out.

1C4E In mid May I had two meetings with Bill Ward. These led to the inclusion of associatve structures as a basic facet of the structural component of the forma.

2 Goals and Constraints

2A The one unifying, pragmatic mandate from the sponsers charges the AHI center with building the AHI Super Screen. It is a large collection of electrical and mechanical devices connected to a computer. The elusive ground rules for construction and guide lines for development have often made progress slow. No one knows what is being augmented, but everyone has firm ideas about features needed in the system. This characterization attempts to isolate and bound a single problem area from the many of the Center.

2A1 The tool is, like many tools, and artifact and an extension of man's basic senses. Not totally unlike paper, pencils, chalk, and blackboards, it is an information storage, retrieval, and manipulation device. Simply, it is a dynamic memory extension.

2A2 (information structure) The tool works on documents, or files. This paper was written exclusively with the current prototype of the eventual tool. The documents are called information structures to emphasize the point that they are blocks of textual and graphical information organized into hierarchical

and associative structures. The total information content is in the structure as well as the blocks, a point many automated textual systems ignore (as do many writers).

2A3 Language, symbolic reasoning, and problem solving also serve as tools [Bru66:37]. The Super Screen might someday have a similar role. It is not inconceivable that the tool could take on currently unknown features that would give its user a power of thought only realizable through the artifact, much as the power of language is reasiable only through speech, and numerical analysis through the computer and its ability to do the computations.

2A4 The unique feature of the Super Screen, over other rudimentary passive tools like paper, is that it is not only responsive but can stimulate the user while he works. In [Bru67:56] the analysis of human implement systems discusses motor, sensory, and ratiocinative capacity amplifers. The threefold amplification is found in the Super Screen.

2A4A It makes huge volumes of information availabe for viewing and manipulation through minute keystrokes.

2A4B In a way that is becoming common with computer-controled devices, the Super Screen reacts instaneously to user commands. The commands, moreover, could be given in structurally complex languages, something impossible with simpler controlling

devices.

2A4C Unlike any previous devices, the Super Screen adds a new dimension to ratiocinative implements, for it observes the operation of the user, his information structures, and his habits. These observations may then be used to provoke thought and stimulate the use of different structures.

2B Emphasis has been placed on the personal nature of the tool not from disregard of the problem of communication among fellow workers but out of caution. The development thus far could easily assimilate many arears of computer science. Further constraints, necessary to bring the problems down to practical levels, are best set by conjunctive disassociation rather than extension of the original premise, especially since the constraints are necessarily artificial.

2B1 Unfortunately, we must not attempt to radically change, through research and development, already existing cognitive tools of the Super Screen user.

2B1A Increased efficency in personal habits may readily be expected of a Super Screen user, just as increased typing speed could be expected of one using a typewriter.

2B1B [Bru66:56] states that "any implement system to be effective must produce an appropiate internal counterpart, an appropriate skill necessary for organizing sensorimotor acts,

for organizing percepts, and for organizing our thoughts in a way that matches them to the requirements of implement systems."

2B1B1 However, we consider it outside the scope of this project to research and make significant contributions toward accelerating the evolution of these skills.

2B1B2 The tool is designed for use in tasks so ill-defined that there are no measures of success or concrete working rules. No one knows of universal rules for taking notes or wrting papers; finding such rules is a project of different orientation.

2B1C Neither are we modeling the person or his thoughts. The aim is to find a mutually agreeable class of models which the user and the computer can easily work on in unison. The ultimate desire is that the information structures enhance the memory of the user beyond the capacity of more conventional artifacts.

2B2 The most difficult criterion to set is an idea of usefulness. The more a feature or aspect of the tool amplifies or stimulates the cognitive or organizational capacity of the user, the more useful it is. A measure of usefullness is untimately a function of the things being augmented. This is discussed in more

detail after the nature of the tool itself is described.

2B2A This definition is more restrictive than the one generally pursued in the AHI Center, where smoothness of the system and application to specific problems often takes priority.

2B2B This does not leave features open to testing. If usefulness is to ever have a well-defined measure, it must be related to a strong, restricted thesis.

2B3 Finally, we make the strong requirement that the services offered to a Super Screen user be limited to state-of-the-art computer science. This avoids the great hangups of natural language and artifical intelligence.

3 Forma and Model

3A The computer and the Super Screen user converse in an artifical language about specific information structures. For the user to realize the full potential of the tool he must have abstract notions of how the structures are built, connected, and manipulated.

3A1 (forma) This set of rules, from which the user can derive specific models of the structures in the computer, is called the forma of the system. The "forma" is in no way related to the programming techniques used within the software of the system. It

is a formalism used to characterize all possible allowable information structures.

3A2 The idea of a forma is similar to two more common notions.

3A2A An analogy can be drawn with logic, where axiom schema are given from which all possible formulas in the system are derivable.

3A2B The forma is loosely related to the "schema" of Piaget. [Fla63:54] considers a schema to be "a kind of concept, category, or underlying strategy which subsumes a whole collection of distinct but similar action sequences." Thus the forma is both a mathematical characterization of possible structures and a general knowledge of how the structures are built and manipulated. The forma should properly include the user heuristics for the problem at hand, but they are deliberately left out.

3A3 A forma has three components:

3A3A (lexical) The lexical component of an information structure is the content of the nodes of the structure.

3A3A1 In the prototype systems, the structures are blocks of text and pictures combined in a graph structure. Each numbered string of text in this paper is such a block. The

rules which describe the possible configurations of text and lines in such a block are the lexical component of the forma.

3A3A2 More precisely, any node in the current system may be a string of characters. The character set has 96 common characters and the maximum length of any such string is 3000 characters. Line drawing is just being implemented and its rules are much more restrictive.

3A3A3 Factors which limit the usefulness of the prototype systems are rarely concerned with the lexical component. Occasionally the need arises for a special character or font, but in general this has little bearing on forma design.

3A3B (structural) The structural component of a forma is the nature of the connections and relations that can be imposed on the nodes of the lexical compontent. [Wh65:215] suggests that the structure of adult mental processes has an associative, fast-acting level and a slower-acting, information-processing cognitive level. If we take (as in [Bru66:48]) the hypothesis that "experience is organized to correspond in some measure to the structure of language," a forma must have both hierarchical and associative aspects.

3A3B1 (hierarchical) A structure is a means of implying

information missing in the nodes of the structure. When the information is supplied purely through the structure, and is independent of the content of the nodes, it is a hierarchical structure.

3A3B1A The tree structure or outline form of this paper is a hierarchical structure. This widely used format has a curious relation to the concrete operation of Piaget [Fla63:164-196]. If information is categorized or put into class heirarchies, any of the groupments and operations can be handled within the tree even though it is not as general as a directed graph.

3A3B1B The utility of the tree suggests that it is adequate for currently conceivable formas. It may even be too general. Natural language is highly restrictive in its structure. The hypothesis mentioned above implies that is the case; one of the experiments outlined below is designed to test it.

3A3B2 (associative) When the information supplied by the structure is simply the fact that nodes are related because of common content, the structure is associative.

3A3B2A All of the attempts to permit the user of prototype systems the facility of associative structures seem to miss something. Techniques such as naming nodes

(putting a word in parentheses at the front) and embedding the names as references in other statements coupled with commands to move the viewing window to a node by pointing at its name have been the extent of user aids for associative structures.

3A3B2B This work has neglected two important aspects of the associative level of mental processes.

3A3B2B1 [Wh65:189-194] finds that associations are made by subjects and lost for further use in fractions of seconds. This means that, if the associations are to be communicated to the computer, the process must be instanteous. Typing in names takes seconds, not fractions of seconds.

3A3B2B2 Secondly, nodes are associatively related if they have a common concept. Previous work within the Center cast associations and linkages in the framework of reference and referent (pointer and name). A more realistic associative scheme would link nodes through common terms or concepts, on the order of inverted files used in information retrieval.

3A3C (dynamic) The dynamic component is the nature of the way things can be changed. It describes the vast variety of ways the text, pictures, hierarchies, and associations can be

manipulated. This aspect of the forma is the abstract way the abstract forma is manipulated. It is not the languages or commands used for the actual manipulations of information structures.

3A3C1 The trend in prototype systems has been to include everything one can think of so that no user will ever feel constrained.

3A3C2 When this is done, however, care must be take to make consistent definitions and conventions. Obscure manipulative features often lead to misunderstandings of the system.

3A4 The most important characteristic of a forma is that the user understand it. An exceedingly complicated scheme which no one can use is not a tool; it is burden. The simplicity of the formas of the prototype systems is one of their best virtues.

3B (model) It seems at times like a trivial point, but a useful distinction can be made between the information structure in the coumpter, the way a person thinks of the structure, and what he sees as he looks at the Super Screen.

3B1 A model is a specific representation of a specific document or information structure. Thus, "information structure" refers to the actual information, as it is stored in the computer. "Model,"

on the other hand, refers to the way the user thinks of his structure in terms of the forma of the system.

3B2 The actual, physical representation of the model is on the Super Screen and is different from the model and the information structure. The user sees characters laid out in a square format; he thinks of the information as being in a tree; and it is really stored using a fancy free-storage technique.

3C (special purpose) In all of the prototype systems there have been a great many special purpose features. These have included the ability to scan every node for textual content and perform arithmetic on lists of specially tagged numbers, and special formatting conventions for tables. Schemes have been proposed for allowing the user to view the structure implicit in highly structured data, in both the lexical and the structural component.

3C1 The features are difficult to discuss relative to the forma. They are aids more closely aligned to motor and sensory capacity amplifiers than to ratiocinative amplifiers.

3C2 They certainly make the systems smoother to use, and they have been invaluable in well defined tasks. They are, however, ad hoc, and have nothing to do with the problem area attacked by the notion of formas.

4 The Communication Language

4A The user communicates with the machine about models by using an artificial dialogue language. This is not the natural language of the lexical component of the forma. It is the language used to direct the machine in its editing, organizational, and retrieval tasks as well as the language used to inform the machine about structural relations and associations.

4A1 Obviously the nature of the language will dictate certain user habits. How this affects the capacity amplification of the system cannot be understood until the amplification itself is determined.

4A2 The dialogue language is nonreflexive. The user currently strikes keys, pushes buttons, and rolls a cursor-positioning device with his hand. The computer responds via the CRT.

4A3 This language plays a vital part in the system. Its flexibility and power crucially determine the success of system features. Stealing the linguistic approach to language analysis and perverting the terms "phoneme," "morpheme," and "structure" leads to rich insights about the contemporary dialogue language.

4B (phoneme) The phonetics of the dialogue language is the set of basic units of communication between the user and the machine, and vice versa.

4B1 There is a strange imbalance between the dialogue language

and natural languages.

4B1A There are about 100 keys in the user's half of the dialogue language. This is more than twice the number of phonemes in English, yet the artificial language is much more restricted.

4B1B This difference would seem reasonable if the communication speed using keys were faster then speech, but it isn't.

4B2 The dialogue language is reflexive only on the phoneme level. The characters displayed on the screen and the ones struck on the keyboard are the same.

4C (morpheme) The morphemes of the dialogue language are the basic units of discourse about the models. The morphemes need not correspond to the atomic units in the structural, lexical, or dynamic components of the forma.

4C1 This is a very important point, for it means that models can exist separately from the language used to manipulate them. In practice this is true even though it has not been previously recognized. When someone is taught about the system, terminology is used which is not in the dialogue language. Users continue to think about the models with these concepts, yet they are not

realized in the language.

4C2 Awareness of this means that the morphemes can vary without changing the formas. Historically, the morphemes of the dialogue language have been considered the forma. The results of designing without consideration of this feature can be seen in the prototype systems.

4C3 Two striking observations can be made when the user's dialogue language is compared to natural language.

4C3A The morphemes are almost identical with the phonemes. This explains the large number of phonemes. But the utility of this feature, even if it is only speed, is still untested.

4C3B The lack of symbolism or abstractness is a glaring deficiency in the language. All the words are concrete. They are always in reference to a specific action about a specific entity in the model.

4D (syntax) The structural or syntactic analysis of the dialogue corresponds precisely to similar analyses of natural language.

4D1 I know of no structurally complex man-machine interactive language. From job-control languages to highly interactive control languages, the structure is always that of a simple linear grammar. The so-called high-level languages of programming

do not apply here. They are used to describe algorithms to a computer program specifically designed to translate the algorithm to another language. These language are not interactive; writing in them takes a good deal of time.

4D1A Except for a few possible queries about the computer's status, the dialogue language is strictly imperative. It always tells the machine to do something, to change the portion under view, to modify a parameter, etc. This could explain the sufficiency of simple structure, for even in natural language imperative statements tend to be short and structurally simple.

4D1B A more adequate explanation seems to be that structural complexity is closely connected to semantic complexity. Simple concrete semantics, such as a forma, do not need structurally complex languages to describe the allowable operations.

4D2 Just as with morphemes, the syntax of the dialogue language has been confused with the dynamic component of the forma.

4E The common analysis methods applied to the study of speed and redundancy of natural language do not apply to the dialogue language.

4E1 In all of the prototype systems, the semantic complexity of the forma is so trivial that single, short commands have been used to implement the user's entire command language. With the current system, the number of commands is becoming larger than the number

of available keys. We do not, however, have anywhere to look for guidance in studing the problem. All work oriented toward the speed and redundancy of artifical language tends to be concerned with communication on the bit and hardware level.

4E2 The computer, through the display screen, can transmit information so fast that the problems again shift back to the user and his ability to read and follow the computer. Observations about response speed in time-sharing systems do not apply to the theoretical level of this discussion. The output channels from the computer handle information much faster than people can assimilate it. The slowness of the input channel is mainly due to the inadequate dialogue language. The computer can certainly handle informaton faster than people, and the limitation on input speed is neither human dexterity nor hardware capability but rather the capability of the control language..

5 Augmentation

5A Through the notions of forma and dialogue language we have attempted to bring the concept of the Super Screen into the arena of rational discussion. The next, and more difficult, problem concerns the ultimate goal--augmentation of intellectual processes. We do not solve this problem; we only offer notes for further thought.

5B The Super Screen, viewed as a capacity amplifier, augments only one portion of the user's intellect--his performance in the task of

information organization on a personal memory level.

5B1 The problem, then, reduces to the overwhelming task of discovering a theory behind the processes of the user as he does the organization, and relating these processes to specific features of the forma or the dialogue language. With such a relationship, we could make hypotheses and test them.

5B2 The notion of performance must not be restricted to the current user task. The availability of the Super Screen may well mean that the user could undertake organizational task which would otherwise be prohibitively difficult.

5B3 The user processes must be restricted if a theory is to be found. This can be done in two ways.

5B3A The first is the task the user is performing. This would be things such as writing a specific kind of paper or computer program, or taking notes while reading a book.

5B3B The other restriction defines a motivation of the user. There are three motivations that appear to lend them themselves to theory.

5B3B1 The Super Screen user could be using the system as a means of personal recall. NLS becomes an extension of his memory. Within it he can store facts, relations among the

facts, and structural information about both the facts and the relations.

5B3B2 The Super Screen can be used as a device to communicate with other Super Screen users. Entirely new approaches can be taken to paper organization if reading will always be through the system. If the intent of the user is communication, the structural devices are standard enough to hypothesize on, and maybe even to test.

5B3B3 Finally, the user may be attempting to correlate information. The process of forming associations, relations, and concepts has been highly studied. If users were set to this task, various current psychological theories could be used as the basis of test concerning both the forma and the dialogue language.

5B4 A serious problem is that individual user ability, independent of the Super Screen, overshadows the results of being augmented.

5C There is a tendency to view the whole project as the construction of a big text editor.

5C1 This is done both by outsiders when the program is explained to them and implicity by insiders when they think about expansion

and new user features.

5C2 This seems to happen for two reasons:

5C2A The text-editing problem is well defined. With simple manipulation of text, graphs, and structure the features and use are already laid out. Authors and editors feel that they know what they want.

5C2B It is very easy to get a big gain in paper production with such a system. Dozens of rough draft copies, all neatly formated, right justified, numbered, etc. accumulate in piles on the desk of the system users. This is fun, but it is also dangerous. The increase in output is often viewed as an intellectual or quality-performance gain--this is a big mistake.

5C3 It is becoming clear that text editing, even with the Super Screen, augments intellect in only a trivial way [Bru63:298-301]. The notions of the Whorfian and Neo-Whorfian [Eng62:24] hypotheses are just not enough.

6 Research Strategy

6A This has little to do with psycholinguistice; it is deliberately left out of the term paper. (It is also a cop-out because I did

finish it.)

6B introspective development

6C technical development

6D empirical research

7 Two Experiments

7A The Value of Restricted Hierarchies

7A1 This is a set of experiments instead of just one. The purpose is to study the effect of structure, as used in NLS, on the process of communication.

7A2 The general idea would be to take a paper or article, and transcribe in into the system in various versions with a variety of structures encompassing a spectrum of complexity.

7A2A The simplest format is the standard paragraph format.

7A2B A next step could be multi-sentence structure, breaking the paragraphs into related units, each unit composed of one or more sentences.

7A2C A final form might have at most a sentence, and

complicated sentences might be broken down further.

7A3 Various hypotheses could then be made about the way users prefer to read using the Super Screen. These could be tested by observing the commands used to move the viewing window over the parts of the document when subjects are asked to figure out what it says.

7A4 Another set of hypotheses could be made concerning the best way to structure the document to convey its contents to a reader [Ronco:19-21]. Subjects could then be asked to read the file, and use the system to make notes. We would test for comprehension and rentention of material.

7A4A The idea that the minimum structure necessary for good comprehension and rentention is a candidate for testing.

7A4B Another candidate is the idea that complicated structures can convey considerable information, provided the user is aware of the meaning of the structure. Here, various structures could be explained before the subjects read the file. Only selected structures would actually be used in the file. The results of the test afterward could give clues abut the utility of the forma in its present form.

7B Kinds of Concepts in Associations

7B1 The motive behind this experiment is to get some clue, however small, about the use of associations in concept formations, the speed at which they occur, how they are eventually used in hierarchicaly concept formation, and what kinds of features could be added to the Super Screen to allow users to record them and retrieve them.

7B2 The experiment is basically clinical. The idea is merely to interview children about a specific topic. I recognize all the problems connected with children, but I see no other place to start. The interviews would merely probe the child for all the information he has about a small number of highly related objects. The hope is that close analysis of the transcribed interviews will lead to ideas about indexing, associations, and the transformation of associations into structural concepts.

7B2A The area of personal relatives is well defined. Questions concerning the location of individuls, such as grandmother, and of sets, such as mother's parents, might initially be answered thourgh associations. As the questioning procedes, and the child's thoughts stabilize, the associations may progress to structures. To observe such a change and to see the origins of a structure could lend a lot to further refinement of the forma.

8 Summary

8A I feel this paper has made progress if two very simple things have been accomplished.

8A1 If the idea of the Super Screen has become firm enough to recieve meangingful criticism. In the past there have only been three kinds of comments: approval through head nodding, suggestions for generally kludgey special purpose features, and the advice that we are all wasting our time. This quickly becomes discouraging. I want to make the global concepts we work with clear enough to be understood and critized by others.

8A2 I also want the global view to be extendable to research strategy and experimental verification. The experiments must directly relate to the forma, dialogue language, and augmentation process if they are to mean anything.

8B I realize that the discussion of augmentation is lacking. It needs much more thought, and I need a lot more reading. However, I also feel that the experiments outlined above could be started. The results would assist in solidifying the ideas about augmentation as well as clearing up more points about the forma and dialogue language by putting them into pratise.

9 References

9A [Bru66] Jerome S. Bruner, Rose R. Olver, Patricia M Greenfield,

et al, "Studies in Cognitive Growth," John Wiley & Sons, Inc., 1966.

9B [Eng62] D. C. Engelbart, "Augmenting Human Intellect: A Conceptual Framework," Summary Report, Contract AF 49(638)-1024, SRI Project 3578, Stanford Recearch Institute, Menlo Park, California (October 1962).

9C [Fla63] John H. Flavell, "The Developmental Psychology of Jean Piaget," D. Van Nostrand Co., Inc., 1963.

9D [Ronco] Paul G. Ronco, et al, "Characteristics of Technical Reports that affect Reader Behavior: A Review of the Literature," Institute for Psychological Research, Tufts University, Medford Mass.

9E [Wh65] Sheldon H. White, "Evidence for a Hierarchical Arrangement of Learning Processes," in "Advances in Child Development and Behavior" Vol. 2, pp. 187-220, ed. by Lipsitt & Spiker, Academic Press, 1965.