SWIM[1]: A "Natural" Interface for the Scientifically Minded Language Learner*

M. Zock, A. Laroui & G. Francopoulo

LIMSI, BP 30, 91406 Orsay, France

Abstract: We describe a system under development, whose goal is to provide a "natural" environment for students learning to produce sentences in French. The learning objective is personal pronouns, the method is inductive (learning through exploration). Input of the learning component are conceptual structures (meanings) and the corresponding linguistic forms (sentences), its outputs are rules characterizing these data. The learning is dialogue based, that is to say, the student may ask certain kinds of questions such as: *How does one say* ⟨idea⟩?, *Can one say* ⟨linguistic form⟩?, *Why does one say* ⟨linguistic form⟩?, and the system answers them.

By integrating the student into the process, that is, by encouraging him to build and explore a search space we hope to enhance not only his learning efficiency (what and how to learn), but also our understanding of the underlying processes. By analyzing the trace of the dialogue (what questions have been asked at what moment), we may infer the strategies a student put to use.

Although the system covers far more than what is discussed here, we will restrict our discussion to a small subset of grammar, personal pronouns, which are known to be a notorious problem both in first and second language learning.

Key Words: Natural language generation, computer-assisted learning, learning by exploration, personal pronouns, problem-solving.

M. Zock received his 1980 PhD in experimental psychology (psycholinguistics) at the University of Paris-Vincennes. Since 1986 he has worked on an ESPRIT project (PALABRE) and an exchange programme between France and Canada for integrated software. He now works in the Language and Cognition group at the LIMSI, Orsay, and organized with Gérard Sabah, the First European Workshop on Language Generation at Royaumont.

His major research interest is the building of psychologically motivated tools to assist the teaching and learning of the tasks involved in text generation.

His major publication (co-edited with G. Sabah) is Advances in Natural Language Generation: An Interdisciplinary Perspective, *vol. 1 (Ablex & Pinter, 1988).*

The Problem of Learning Language: Language Learning as Problem-solving

Speaking, reading and writing are skills which have to be learned. These skills can be learned in many ways: by imitation, by explanation, by exploration, etc. We will be concerned with the latter approach, i.e. the inductive language learning.

Language learning can be viewed as a special case of problem solving in which the learner tries to build and intelligently explore a hypothetical search space.[2] If this view is correct, then two sets of questions arise immediately. On one hand we wish to know:

- what the nature of this search space is (what are the variables?);
- how it is built (incremental learning: local vs global view);
- how it is explored (strategies: intelligent, opportunistic vs. systematic search).

On the other hand, we wish to investigate how:

- the knowledge at the outset and
- the ordering of the data will affect building and searching of space.

Typically we do not learn from scratch, nor is it

Computers and the Humanities **23**: 411–422, 1989.

likely that we encounter either well-ordered data, or a complete set of examples. Natural learning is incremental.

Obviously, these facts have consequences. They imply that:

- initial knowledge, in particular, knowledge of other languages may bias the kind of variables (attributes or hypotheses) considered, i.e., included in the search space.
- the nature and the order of the data, (the examples encountered by the student and their sequence) may determine what rules are likely to be inferred at what moment, and finally
- rules are inferred from incomplete data (incremental learning). Furthermore, the same data may be characterized in different ways. That is, several equivalent descriptions may be inferred from the same data set. Which of these descriptions turns out to be the most adequate generally cannot be established until one knows the complete data set. Thus, rules may have to be revised in the light of new evidence. Consequently, errors are not only unavoidable parts of the learning process, but also an indispensable source of information for the learner.

The Problem of Learning How to Learn

As we have seen, learning can be seen as searching. Actually, teaching, as well as learning, can be conceived of as problem solving or reasoning in an information-exchange environment. There is a sender, a goal, a message and a receiver. The *sender* may be a native speaker, a teacher, a parent, a book or a computer. The *goal* is the task or final performance (output). In our case it is knowledge of how to produce sentences in French. The *message* is the input to the learning component: examples from which the rules have to be inferred.[3] The *receiver* or learner can be any system, natural or artificial, capable of perceiving, memorizing and analyzing a set of data and drawing the necessary conclusions: a child, a student, or a computer program.[4]

Learning occurs in various settings. Depending on the order of the examples and the control of the information-flow, we speak of natural-, experimental-, or institutional-settings.

Natural learning is characterized by the fact that there is no clearly defined learning objective,[5] by noisy and heterogeneous material, and by unordered examples. The underlying regularities are thus multiple, diffuse, and hard to perceive. Experimental-learning and teaching, on the other hand, have a learning objective, the material is error-free, homogeneous and coherently ordered according to some point of view (learner or teacher).

Whereas *experimental learning* can be characterized by the following sequence: (a) encountering the data, (b) analysis, (c) building and testing of hypothesis, (d) feedback and (e) proof or demonstration of the theory, *traditional teaching* goes generally through the following stages: (a) exposition, (b) practice, (c) testing and (d) evaluation. This can be schematized as follows:

Teacher: sets the task and presents the learning material;
Student: analyzes the data;

Teacher: provides a set of examples;
Student: practices;

Teacher: asks questions to test the gained knowledge;
Student: answers the questions;

Teacher: evaluates the answers, provides feedback (explanations) and organizes future data as a function of actual performance;
Student: integrates the feedback into the knowledge base and corrects misconceptions;

As one can see, the information-flow here is entirely teacher-controlled. He is the one who sets the task, provides the examples and the feedback. Consequently, the teacher decides the nature and the order of the material to be learned.

There are two major shortcomings in this approach. Not knowing what information is needed by the learner, the teacher may present the wrong data, i.e. data which are inappropriate, hence misleading. More importantly, the student is only loosely integrated in the learning process. Instead of being active, generating and testing plausible hypotheses (discovery learning), he primarily reacts to questions. Thus, it may happen that the student perceives his task as the learning of the material (rote learning) rather than the underlying principles.

Ignorance of what or how to learn may result in (a) learning the unintended, (b) poor problem-solving skills, or (c) little transfer. As long as the learner does not go *beyond* the *information given* (the concrete word level), he cannot transfer the gained knowledge to similar situations, because the perception of similarity presupposes abstraction.

Given these criticisms, it would be useful to have a system which has the qualities mentioned above without having the drawbacks. A good learning environment should be both flexible and constraining enough:

- to allow for simulation of real communication, that is to say, to provide a setting where both participants can take the initiative and control the information-flow,
- to ensure the learning of the appropriate material (i.e., *what to learn*) as well as the necessary problem-solving skills (the methods, i.e., *how to learn*).

A computer program could be such an environment. It would offer different kinds of information (see below: communication mode) while answering the student's questions as he goes along generating and testing different sorts of hypotheses.

The Cognitive Engineer's Task: To Provide the User with a Friendly Interface

We will describe here a system under development, whose major goals are:

- to provide an environment which allows communication between a learner (student) and an expert (in our case the system);
- to simulate the information-processing aspect of natural learning, i.e., the inductive learning of grammar rules to generate sentences in French.
- to allow teachers and psychologists to test various theories.

The system we are developing is designed to help the student build the search space, i.e. the set of all attribute-value pairs. The learner has to discover how to explore it. By applying a given set of operators and by watching the outcome, he can test (a) which information is relevant, and (b) to what it is relevant (to syntax or morphology). However, in this kind of dialogue (controlled trial and error) the system not only answers the questions asked by the learner, but also assists him to determine what questions are meaningful in this context.

Learning, be it by man or machine, implies exchange of information between two systems, for example, a native speaker (expert) and a foreigner (learner). We will start by describing some of the features a system needs to have in order to allow for such an information exchange. We will then give a detailed example, showing what such a dialogue between a human learner and the machine might look like. Finally, we will discuss whether machines can acquire linguistic competency in a humanlike way.

Before showing how the system works, let us specify more clearly the learning objective.

The Student's Learning Objective

The learner's task consists of incrementally learning the morpho-syntactic rules of personal pronouns in French. More precisely, the student is expected to acquire the necessary knowledge in order to generate sentences composed of several pronouns (see examples in Figure 2).

In order to achieve this goal, he has to learn:

- how to express a given concept (morphemes);
- how to linearize these concepts (sentence patterns), and
- under what conditions (rules) to use each of these words or sentence forms.

As one can see from the data, pronoun-constructions in French can be quite complex.[6] This complexity is due to:

- the *number of features* necessary to determine syntax or morphology;

Concepts	Expressions	Rules for choosing among the forms (morphology)	
SPEAKER	je, me, moi, - nous	if	*Syntactic function* : direct object
LISTENER :	tu, te, toi, - vous	&	*Person* . third
ELSE :	il, elle, ils, elles	&	*Reflexive* . no
	le, la, les, lui, leur	&	*Quantity* : definite
	on, en, se, soi, eux	&	*Number* : singular
		&	*Gender* : female
		then	*Direct Object* --> la

Figure 1.

SYNTAX: (word order patterns)		
(a) S-DO-IO-V	Je la lui présente.	I introduce her to him.
(b) S-IO-DO-V	Je te la présente.	I introduce her to you.
(c) S-DO-V-prep-IO	Je te présente à elle.	I introduce you to her.
(d) S-IO-V-prep-IO	Je lui parlerai de toi	I will tell her about you.
(e) V-DO-IO	Présente-la moi !	Introduce her to me.
(f) neg-DO-IO-V-neg	Ne la lui présente pas !	Don't introduce her to him
(g) neg-IO-DO-V-neg	Ne me la présente pas !	Don't introduce her to me.
(h) neg-DO-V-neg-prep-IO	Ne me présente pas à elle !	Don't introduce me to her.
(i) neg-IO-V-neg-prep-IO	Ne lui parle pas de moi !	Don't tell her about me.
S: subject, DO: direct object, IO: indirect object, prep: preposition, neg: negation, V: verb		

Figure 2.

Part of speech : (noun, pronoun)
Je parle à *Pierre*. (noun)
Je *lui* parle. (pronoun)

Syntactic function : (subject, direct object, indirect object)
Il écrit à Pierre. (subject)
Paul *lui* écrit. (indirect object)

Sentence-type : (declarative, interrogative, command)
Tu *me* le donnes? (interrogative)
Donne-le *moi* ! (command)

Polarity : (positive, negative)
Donne-le *moi* ! (positive)
Ne *me* le donne pas! (negative)

Communicative roles : (I, you, he)
Je te *le* donne. (indirect objet : you)
Je *le* lui donne. (indirect object : he)

Number : (singular, plural, indefinite)
Je te *le* garde. (singular)
Je te *les* garde. (plural)
Je t' *en* garde. (indefinite)

Gender : (male, female)
Je *le* vois. (male)
Je *la* vois. (female)

Verbconstruction : (type of complement (direct vs. indirect object), type of preposition, reflexivity)
Je vois Marie. --> Je *la* vois. (direct object)
Je parle à Marie. --> Je *lui* parle. (indirect object)

Semantic features : (animate, inanimate)
Il m'emmène à Rome. --> il m'y emmène
Il me présente à sa mère. --> il me présente à *elle*

Figure 3.

- the *structure of these features*: if one compares (a) and (c), one will notice that the form of the indirect object (lui vs elle) depends on the value of the direct object (cross dependency);
- the *interdependance of syntax and morphology*: practically all variables, except *number* and *gender* are relevant both for syntax and morphology. Futhermore, the position of the direct object pronoun may depend on the value of the indirect object (compare (a) and (b) here above). In other words, changes of morphology may cause changes in syntactic structure.
- the *various knowledge sources*: the determination of morphology and syntax requires information about the *referent* (number, gender, animacy), *text functions* (syntactic status of noun-phrase: noun vs. pronoun, topicalisation, person), *polarity* (positive/negative), *speech-act* (statement/question/command), *verb-construction* (type of complement: direct/indirect, type of preposition: à, de), etc.

Given these intricacies it is easy to understand why students so often fail to learn these rules. Assisting their learning, or modelling the underlying process is thus a challenging task.

How to Integrate the Learner into the Process?

If one accepts our view of learning, then the problem of the student is to find out how to build and how to reduce intelligently the search space. The system will help the student in various ways.

First of all, it will answer certain kinds of questions:

(a) How does one say: ⟨idea⟩?
(b) What would happen if: ⟨syntactic operation⟩?
(c) Can one say: ⟨linguistic form⟩?
(d) How should one say: ⟨linguistic form⟩?

All these questions occur in some form or another in natural settings. The following examples will illustrate these strategies or testing modes:

Question	Answer
How does one say ? parler (Paul, Marie) talk (Paul, Mary)	Paul parle à Marie (Paul talks to Mary)
What would happen, if "Mary" were pronominalized ?	Paul lui parle
Can one say: "je lui pense"?	No.
Instead of saying. "je lui pense", *how should one say* ?	This depends on what you want to say: [1]
	Je pense à elle. (I think of her.) Je pense à lui. (I think of him.) Je le pense. (That's what I mean.)
Why does one say : "Je le pense"	Explanation given by the system.

Figure 4.

These strategies are complementary in that they correspond to different learning needs. They provide different kinds of feedback. The first two methods (the inductive approach) seem useful if one does not have much knowledge yet. The third one allows to test the degree of generality or the extension of a given rule (deductive reasoning), the fourth method provides additional information in case of incorrect performance, while the last question may either confirm a hypothesis, or correct a misconception.

Second, the system should show how to reach the solution (the demonstrative mode). This might be helpful if the student gets stranded, not knowing what to do. In this case the system takes over, showing how information may be processed. By watching the system, the student may learn how to explore, i.e., how to generate and test a set of hypotheses.

Third, the system keeps a record of the whole dialogue. Such a trace has several advantages: it allows the student to verify, to explain and to remember. He may thus (a) check the consistency of the rules, (b) justify a given conclusion in the light of evidence and (c) reorganize his knowledge base. This last possibility should enhance his perception of underlying regularities.

Psychologists could use this trace to infer the student's learning strategies. The rules a student has been testing at a given moment may be inferred on the basis of the nature and order of the questions being asked.

Finally, teachers could use the trace-function to gain feedback concerning the optimal order of presenting the data. By varying the nature and order of information, they can determine experimentally the complexity of the data (examples, rules), and thereby the relative efficiency of teaching-strategies.

Description of the System[7]

The program works interactively. The user is given a set of options (Figure 5) from which he has to choose. The system converts this input into a conceptual graph (Sowa 1984) and computes the adequate output, i.e. linguistic form. Input are meanings, or conceptual structures (what to say), output are sentences (how to say).

The dialogue is initiated by giving the system the communication mode. Assume that the mode is "How does one say ⟨idea⟩?" In that case the system displays the following menu:

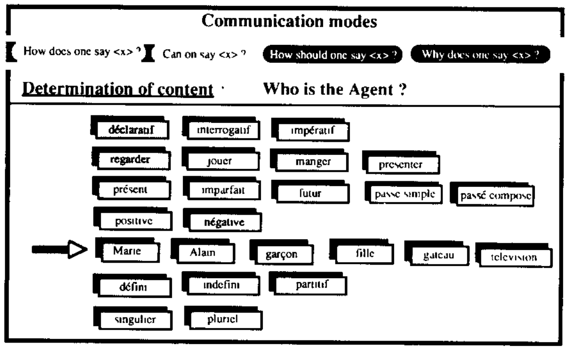

Figure 5.

By choosing among variables like:

sentence mode: déclaratif, interrogatif, impératif

type of verb: manger, regarder, jouer, présenter

tense: présent, imparfait, futur, passé composé, passé simple,

sentence form: positive, négative

discourse objects: garçon, fille, gateau, télévision, Alain, Marie[8]

definiteness: défini, indéfini, partitif

number: singulier, pluriel

the user tells the system *what he wants to say*. As the dialogue develops, the system incrementally builds the underlying meaning and outputs the corresponding form.

Past this point, the user has various options:

- either he builds a completely different sentence, in which case he would have to go through the whole routine as depicted in figure 5, or,
- he just changes the value of one of the parameters, i.e. variable. For example, he can change the verb, the tense, the sentence mode, etc. To do so, he clicks the particular value of the variable and specifies the new value. The system will change the meaning accordingly; it replaces the old value by the new one and outputs the new sentence form. For example, if the reference form is:

(a) Le garçon regarde *la* fille (the boy watches the girl)

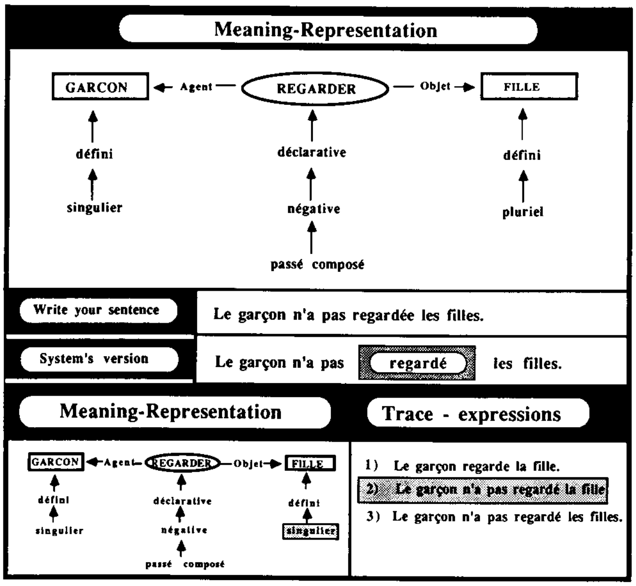

Figure 6.

and if one asks the system to replace the feature "singular" by "plural", it produces:

(b) Le garçon regarde *les* filles (the boy watches the girls)

As one can see from figure 6, the screen is divided into three parts. The large window contains the sentence under construction (its meaning representation and corresponding form) and two smaller windows (bottom) which represent the user's *memory of meaning* and *memory of form*. The latter, called "trace of expression," is a data-base, containing all the sentences encountered so far. This base can be organized (manually or automatically) according to the user's needs. For example, all sentences in interrogative-negative form may be grouped together, irrespective of the order in which they were constructed. The memory of meaning window contains the sentence's underlying conceptual representation. It is activated by clicking any of the forms contained in the data base.

The idea behind this separation is to allow the user to make a contrastive analysis of meaning and form between two sentences. Choosing a sentence in the "memory of form" window gives a conceptual graph representation of the sentence's meaning in the "memory of meaning" window. By comparing the surface form and the underlying meaning of two sentences, the user can appreciate step by step the relationship between meaning and form. The critical feature, the one that is responsible for the difference of form, is highlighted by the system. In our example it is the value "singular."

One last option is called "transformation." This allows the user to perform certain transformations such as passive voice or pronominalization. In the latter case, the system will ask the user to specify which element he wants pronominalized (one argument or both). Assume that, starting from the conceptual structure underlying sentence (a), the user wants to pronominalize respectively the agent, the object, and finally both arguments. In that case he would get the following outputs:

Il regarde la fille. (he watches the girl)
Le garçon *la* regarde. (the boy watches her)
Il la regarde. (he watches her)

By comparing these sentences with the base form, the student should notice certain differences and draw the necessary conclusions. For example, given the data he may conclude that:

R1: if the direct object is pronominalized, then it moves in front of the verb (syntax).
R2: case (syntactic function) is morphologically relevant:
if the subject is pronominalized its form is "il,"
R3: if the direct object is pronominalized its form is "la".

From now on we are in a loop, the dialogue being basically the same. However, in each cycle the hypothesis to be tested is likely to be different and it is interesting to watch how a student proceeds in order to acquire competency. What does he want to know? Is he systematic? What kind of strategy does he use (breadth first, depth first etc.)? Under what conditions does he change his method?, etc.

The learner's problem is three-fold, he must find out:

- which parameters (attributes) are relevant,
- to what linguistic component they are relevant (syntax and/or morphology), and
- to what extent they are relevant.[9]

A student may thus want to know:

- whether the variable *gender* is morphologically relevant;
- whether this is the only relevant variable, or if other variables come into play;
- whether it is relevant for all cases, irrespective of, for example, communicative role, negation, or sentence mode (compare (e) and (g));

Every time the student is given control, he can choose two things: (a) the kind of information he wants to convey (what to say), and (b) the dialogue-mode, i.e. how does one say ⟨idea⟩?, Can one say ⟨linguistic form⟩?, etc.. The following diagram illustrates the information flow.

This kind of environment has basically three functions:

- to answer different kinds of questions
- to convert meaning into form, and

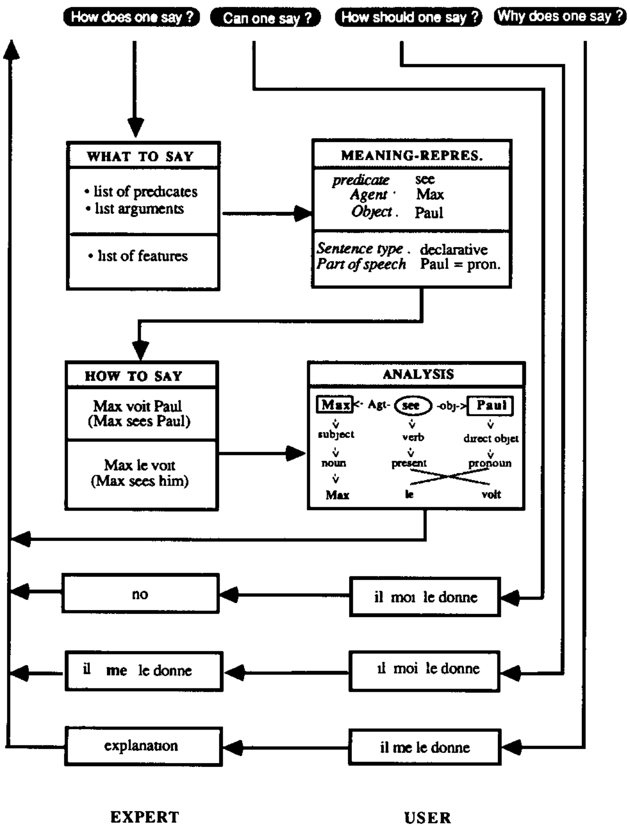

Figure 7.

- to help the student to discover how changes in meaning are reflected in changes of form.

It should be noted, that the student has most of the control. The following examples should give an idea of the dialogue. These hypothetical dialogues serve illustrative purposes. However, we believe that they are reasonably close to what might be encountered in a real experimental session.

Sample Dialogue 1

The dialogue mode is : *How does one say* ⟨idea⟩? Figure 8 contains three columns which express respectively the student's intentions, i.e. what he wants to say, his observations, and his conclusions with respect to syntax and morphology.

Having generated the following proposition:

voir (Max, Paul)
see (Max, Paul)

he wants to know what would happen if *Paul* were pronominalized. The system generates the following answer:

(1) Max le voit

The student analyzes this sentence and draws as conclusions Rule 1 and Rule 2 mentioned here above. It should be noted, that rule 2, although correct, needs further refinement. It is underspecified as it lacks information concerning *number* and *gender.*

During the next cycle the student asks what *would happen if* "Paul" was replaced by "Mary". The system answers:

(2) Max la voit

The student concludes that *gender* is not relevant with regard to word order, but is a necessary condition to determine morphology (Rule 3). This latter kind of knowledge could be expressed as:

R3: if *Part of spech*: pronoun
& *Syntactic function*: direct object
& *Gender*: female
then *Pronoun*: la
else
if *Gender*: male
then *Pronoun*: le

In the next question he is concerned with the

INPUT (conceptual structures)	OUTPUT (sentences)	CONCLUSIONS (rules inferred)
1) see (Max, ***Paul***) Paul = pronoun ->	**Max le voit**	
	DO precedes verb	**Syntax :** *Syntactic category & Syntactic function* of the referent are syntactically relevant
		R1 : if *Syntactic category* · pronoun *Syntactic function* · direct object then pronoun precedes the verb· S-DO-V
	pronoun = le	**Morphology :** *Syntactic function* is morphologically relevant
		R2 if *Syntactic function* direct object then *pronoun* le
2) see (Max, ***Mary***) Mary = pronoun ->	**Max la voit**	
	position of DO consistent with R1	**Syntax :** *Gender* is syntactically not relevant
	change in form	**Morphology :** *Gender* is morphologically relevant
		R3 : if *Syntactic category* pronoun *Syntactic function* · direct object *Gender* : female then *pronoun*: la
		R4 : (refinement of R2) if *Syntactic category* · pronoun *Syntactic fucntion* direct object *Gender* : male then *pronoun*: le
3) see (Max, ***children***) children=pronoun ->	**Max les voit**	
	no change in position	**Syntax :** *Number* is syntactically not relevant
	change in form	**Morphology :** R5 : if *Syntactic fucntion* . direct object *Number* · plural *Gender* : male (*) then *pronoun*: les
4) ***speak*** (Max, ***Mary***) Mary = pronoun ->	**Max lui parle**	
	IO precedes verb (see R1)	**Syntax :** *Case* is syntactically relevant.
		R6 : if *Syntactic category* pronoun then Subject - Indirect Object - Verb
		R7 : Generalization of R1 & R6 if an object is pronominalized then Subject-Object-Verb
	change in form	**Morphology :** *Case* is morphologically relevant (see R2)
		R8 : if *Syntactic function* . indirect object *Gender* · female (**) then *pronoun*. lui
5) speak (Max, ***Paul***) Paul = pronoun ->	**Max lui parle**	
	consistent with all preceding rules	**Syntax :** *Gender* is syntactically not relevant
	no changes in morphology	**Morphology :** The *Gender* of the *indirect object* is morphologically not relevant In consequence, relax the gender constraint of R8
		R9 : if *Syntactic function* indirect object then *pronoun*· lui

(*) Since *gender* was relevant for the *singular*, the student assumes that it is also relevant for the *plural* (overspecification),

(**) Since *gender* was relevant for the *direct object* (R3), the student assumes that it is also relevant for the *indirect object*

Figure 8.

relevancy of *number*. He asks: *what would happen if* the direct object were "*children*" (les enfants)? The system's answer

(3) Max les voit

allows him to conclude that *number* is relevant for morphology but not for syntax, as there are no changes in word order, but there is a change in form. This fact is encoded in the following rule:

R4: if *Syntactic function*: direct object
& *Gender*: male
& *Number*: plural
then *Pronoun*: les

It is interesting to notice that this rule is too specific, because *gender* is not a necessary condition. However, this conclusion is perfectly reasonable given the data encountered so far. *Gender* was a necessary condition for *singular* (see rule 3), and so far there has been no evidence to the contrary. Consequently, the student has no way to conclude from the data that for direct objects *gender* is not relevant for the *plural*. (The only reason we could think of why a student might consider this last hypothesis, would come from his knowledge of another langugage which has this very same property.)

The fact that *gender* is only relevant for the *singular* has procedural implications; namely the attribute *number* should be processed prior to *gender*. The former being more informative than the latter.

In the following cycle (sentence 4) the student changes the proposition altogether, asking the system how one would say:

parler (Max, Paul)
speak to (Max, Paul)

when the indirect object (Paul) is pronominalized. This would yield the following sentence:

(4) Il lui parle

From that he may conclude, that the indirect object precedes the verb (Rule 5). Recognizing the similarity with rule 1, i.e. recognizing the fact that the syntactic status of the object (direct object vs. indirect object) does not affect word order, he may generalize these two rules and replace them by rule 6:

R6: if an object is pronominalized, it precedes the verb

This rule is more general than the former ones, in that the distinction between direct and indirect object has been dropped. It should be noted, however, that this rule, even though correct in the light of evidence, i.e. data encountered so far, is too general. For example, it does not apply to sentences composed of two objects (three place predicates). In other words, this rule needs refinement, in the form of additional constraints.

With respect to morphology, the student concludes that the attribute *case* (syntactic function) is relevant, which yields the following rule:

R7: if *Syntactic function*: indirect object
& *Gender*: female
then *Pronoun*: lui

Again, the morpheme is overspecified, because *gender* is not a necessary condition. Having noticed that *gender* was relevant for direct objects (rule 3) the student overgeneralized, assuming that it was also relevant for the indirect object. It is noteworthy, however, that this particular overgeneralization does not produce incorrect results.

Finally, the student asks the system to replace *Mary* by *Paul*. Getting the same answer as in 4, he concludes that for indirect objects the *gender* is irrelevant for syntax as well as for morphology. Consequently, he relaxes the gender-constraint of rule 8. Once again, this conclusion is valid only with respect to the set of examples encountered.

Can one say ? (user's question)	(system's answer)	Hypothesis (implicit question)	Conclusion (answer)
1) Il me *le* donne.	Yes.	Q1 Do both objects precede the verb ?	*Conclusion* 1 · Yes
		Q2 Which one precedes the other ?	*Conclusion* 2 : The indirect object precedes the direct one : S-IO-DO-V
2) Il me *la* donne.	Yes.	Q3 Is the *gender* of the direct object syntactically relevant ?	*Conclusion* 3 : No.
3) Il me *les* donne	Yes	Q4 Is the *number* of the direct object syntactically relevant ?	*Conclusion* 4 No
4) Il *te* le donne	Yes.	Q5 Is the *person* of the indirect object syntactically relevant ?	*Conclusion* 5 . No
5) Il *lui* le donne.	No, one says : Il le lui donne.		Correction and refinement of questions 2 and 5. The variable person is syntactically relevant.
			Conclusion 6 : if Person-IO third then . S-DO-IO-V
			Conclusion 7 : if Person-IO : first or second then : S-IO-DO-V
6) Il le *se* garde	No, one says · Il se le garde		*Conclusion* 8 (refinement of conclusion 7).
			if Person-IO: third V-construction : reflexive then : S-IO-DO-V
7) Il s'*en* moque.	Yes		confirmation of conclusion 8
			Conclusion 9 If number-Do indefinite then S-IO-DO-V
8) Il *en* lui donne.	No, one says Il lui en donne.	Q6 Which one of the two objects precedes the other ? (see concl.2)	The indirect object precedes the direct one Consequently conclusion 6 has to be refined.
			Conclusion 10 · If Person-IO : third Number-DO · indefinite then S-IO-DO-V
9) Il m'*en* donne.	Yes.		confirms conclusions 7 and 9
			With regard to the examples given in 7 and 9 we may relax the person-constraint of conclusion 10.

Figure 9.

Sample Dialogue 2

This time the dialogue-mode is: *Can one say* ⟨ *linguistic form*⟩? This mode is important in that it allows to receive negative evidence. So far the student has received only correct sentences. However, knowing what forms are incorrect is often necessary to avoid overgeneralizations.

The three columns of figure 9 correspond to the student's questions, his hypotheses, and his conclusions. The controlled variable (a change of attribute, or a change of its value) is highlighted.

The figure being rather self explanatory, we will make only some brief comments. At stage 3 the student wants to know whether the *communicative role* of the indirect object, the attribute *person*, is syntactically relevant. From the data he has seen, he concluded that this was not the case. However, this conclusion, even though correct with respect to the data, has to be revised in the light of new evidence (see sentence 4).

It is interesting to note, that the student would probably never have drawn this conclusion, if sentence 4 had preceded sentence 3. In other words, he would have noticed the relevancy of the attribute *person* immediately, hence avoid conclusion 5.

il *le* lui donne (he gives it to him/her)
il te *le* donne (he gives it to you)

This shows how order of data is a critical variable, determining the efficiency of rule-inference, i.e. what conclusions are drawn at what moment.

Can Machines Acquire Linguistic Competency in a "Human Way"?

In fact there are three questions:

- Can machines learn?
- Can they learn in an intelligent or "human" way?
- What kind of knowledge would a computer

program need to have in order to learn the rules we have been talking about?

The answer to the first question is clearly yes (Michalski, Carbonell and Mitchell, 1983). The latter two questions are more controversial. Let us begin with the last one.

Inductive learning basically consists of drawing conclusions from the similarities and differences of abstract data descriptions (contrastive analysis). The crucial points are thus *data description* and *analysis*:

- in what terms should we characterize the data?
- what additional kind of knowledge is needed to infer the rules?

Obviously, a system capable of performing the kind of learning we have been talking about would have to be able to parse the sentences, that is, it would have to produce as output an adequate description of the input sentences described above.

This raises a terminological problem. Data can be described in various ways. Different descriptions can be functionally equivalent.[10] Clearly, the choice of metalinguistic terminlolgy differs depending on whether the goal is machine learning or modelling "human" learning. In the first case, the problem is descriptive adequacy, whereas in the second case we deal with an additional constraint, that of the universal status of the terminology. Do all humans, irrespective of culture and education, use the terminology of linguists? Is there a universally shared subset of metalinguistic vocabulary? In the absence of answers to these empirical questions we will stick with the terminology currently used in computational linguistics.

A different, but related problem is the question of how a system may be enabled to draw conclusions from a set of data (infering general rules).

As we have said above, generalizations are made on the basis of contrastive analysis. In order to allow for such generalizations, the learning component needs a hierarchically structured metalanguage, that is, a vocabulary whose low level concepts (primitives) are subsumed by more highly ordered, abstract forms of knowledge. For example:

masculine & feminine → *Gender*
singular & plural → *Number*
subject, direct object → *Case*

We will now turn to the question of whether computers can learn in an intelligent or "human" way? Obviously this question raises the problem of what intelligence is. Instead of answering this question, we will focus on two aspects of intelligent learning, namely economy and flexibility of methods.

Exhaustive search is neither natural nor economic. Since memory is associative, we find it hard to be consistently systematic. Like gamblers, we tend to use more or less risky search methods (opportunistic search).

People generally have a set of methods and a separate component (critique) for evaluating these strategies with respect to their relative efficiency. As different problems require different problem-solving methods, it is very unlikely that there is a unique, universal problem-solving method. People tend to be opportunistic in their approach rather than systematic or scientific. Both the nature of strategies and the depth of processing will vary with the needs of the learner. Corollarily, it is equally unlikely that one finds the optimal method immediately, since one operates on incomplete data. Inductive learning is typical incremental. Hence methods have to be adapted or gradually refined in the light of new evidence.

Intelligent learning is thus intimately linked to strategic knowledge[11] and to (more or less) general information-processing principles.

These principles may be expressed in terms of simplicity, informativeness, generality, and so forth.

For example, the notion of *simplicity* may be used to choose among different options. In fact, a learner could hypothesize that two-place predicates (eg. "to see") are easier to process than three-place predicates (eg. "to give").

The notion of *information* is related to efficiency. It can be used to reduce the search space. This claim is substantiated by the fact that rules governing morphology of first and second personal pronouns (I, you) are generally learned faster than those which determine the form of the third person (he). The former carry more information than the latter.

In conclusion, we believe that, in principle, certain aspects of intelligent learning could be modelled by computer. However, before trying to model human learning, it may be worthwhile to start gathering data on how humans learn. This is precisely our goal. By watching how people use this tool, i.e. by keeping a trace of the dialogue, one should be able to infer the strategies they use.

Conclusion

We have described a system under development that is meant to be a tool for theory builders (computational linguists, cognitive psychologists), application designers (language teachers) and end users (students). The system is meant to assist linguists, psychologists, teachers and students in their respective tasks: elaborate grammars, model learning, optimize teaching and learning strategies.

The emphasis in this paper has been on *learning* rather than on *teaching*. For the time being the task of learning is to be performed by a human, however, in principle it is possible to extend the system so as to allow for automatic learning, the ultimate goal being to model human-like behavior (opportunistic, or intelligent search).

Computers with their large, indelible memories are powerful tools. They allow us to control virtually any number of parameters. Consequently, one can trace a reasoning process or test a given theory, i.e. determine empirically how different variables affect the efficiency of learning, and so forth.

This has an interesting consequence with respect to theoretical commitments. Instead of claiming an all-encompassing model or theory, one can write a program general and flexible enough to permit the testing of various theories. That is what we are trying to do.

Watching how people use the tool, we may gain important insights about the way humans learn, and thus eventually move from artificial to natural intelligence.

Notes

* The research here described has been supported by Apple Computers, Europe.

[1] The acronyme SWIM stands for: "See What I Mean?"

[2] A search space may be characterized by the possible combinations of perceived attribute-value pairs. *Attributes* are in our case metalinguistic variables such as "person, number, gender," the corresponding *values* are: first, second, third person, singular, plural, etc.

[3] This message has to be interpreted. Thus the learning task is not the surface form of the message, i.e. words and sentences, but the underlying principles (abstractions: rules and sentence patterns) allowing their generation. While some forms (e.g. words) have to be learned, they generally serve for illustrative purposes. Rote learning of the entire set of surface forms (words and word combinations) is not only inefficient, but in fact impossible, because of time constraints: there are more possible combinations than we have time to learn.

[4] It should noted, however, that we are not dealing here with children learning a first language. Instead we would like to model the inductive approach taken by the scientific minded foreign language learner.

[5] One may object that there is a global goal, namely learning the language. However, it seems to us that the primary goal is communication rather than attaining a local objective like, let us say, learning the pronoun system in French.

[6] One may object that there is a global goal, namely learning the language. However, it seems to us that the primary goal is communication rather than attaining a local objective like, let us say, learning the pronoun system in French.

[7] The current version, written by A. Laroui in Lisp, runs on a Macintosh. To date, only the first communication-mode ("How does one say?") is implemented.

[8] This list is a kind of knowledge base, i.e. a set of facts a potential user may talk about. This base is limited in scale, and arbitrary, in that it is given by the system. Please remember, this system was initially designed only for the learning of personal pronouns. In order to generalize our approach we have to enhance the power of the conceptual component. The important point is that, by feeding nouns and verbs into the knowledge base and by choosing among these entities, the student signals what he wants to say.

[9] This last problem, which consists in finding the right degree of generality (underspecification vs. overgeneralization), is particularly delicate in that conclusions have to be reached on the basis of incomplete data (incremental learning).

[10] This fact is illustrated by the variety of parsers. Parsers analyze sentences and assign them descriptions on various levels such as: part of speech, syntactic function, case-roles and so forth. For a review of the state of the art see King (1983), or Winograd (1983). For a French parser see Francopoulo (1986).

[11] These strategies could either be part of the system, in which case they must be explicit (one needs a model), or they could be part of the learning process, in which case the system learns not only domain-specific knowledge, but also methods of how to learn (metaknowledge).

References

Francopoulo, G. "Machine Learning as a Tool for Building a Deterministic Parser." In *Second Austrian Congress on Artificial Intelligence, GWAI 86*. Ed. C. Rollinger and W. Horn. Springer Verlag, 1986.

King, M. *Parsing Natural Languages*. New York: Academic Press, 1983.

Michalski, R., J. Carbonell and T. Mitchell. *Machine Learning: An Artificial Intelligence Approach*. Palo Alto, CA: Tioga Publishing Company, 1983.

Sowa, J. *Conceptual Structures: Information Processing in Mind and Machine*. Reading, MA: Addison Wesley, 1984.

Winograd, T. *Language as a Cognitive Process*, Reading, MA: Addison Wesley, 1983.

Zock, M., G. Sabah and C. Alviset. "From Structure to Process: Computer Assisted Teaching of Various Strategies for Generating Pronoun-Constructions in French." In *Proceedings of the 11th Coling*. Bonn, 1986.