MONTE CARLO SIMULATION OF LANGUAGE CHANGE
IN TIKOPIA & MAORI*
by

Sheldon Klein, Michael A. Kuppin
& Kirby A. Meives
Computer Sciences Department
1210 W. Dayton Street
University of Wisconsin
Madison, Wisconsin 53706

1.   Introduction

The use of Monte Carlo Simulation with micro socio-
linguistic models permits testing of many hypotheses un-
verifiable by any other known method. The methodology un-
derlying the research described in this paper in outlined,
and, to some extent, justified in [20-22]. Basically, the
technique requires a simulation model with the following
subcomponents:

a) A stochastic socio-demographic model of a speech
community for the starting date of the simulation.
This model governs the conversational interaction
patterns among members.

b) A metamodel of significant historical events and
changes during the simulated time period for use
in generating periodic revisions in the basic model
mentioned above.

c) Individual models of members of the society in the
form of dynamically modifiable inputs to the para-
meters that serve as inputs to the rules of the
basic model. The model of each individual also
includes one or more grammars that may be filled
with generative rules for several languages.

d) A language learning component, both for children
and adults. This module permits the generation and
parsing of sentences using rules from the grammars
of specified members of the simulation. The learn-
ing component makes it possible for a child born
during the simulation to acquire the language or
languages of his speech community through conver-
sational interaction with other members of the so-
ciety, and permits an adult either to modify one
of his grammars in response to some contemporary
linguistic innovation, or to acquire a new lan-
guage with rules stored in a separate list. The
learning component to be used in the system is a
greatly improved version of the AUTOLING system
[22 ,23].

A preliminary testing of the simulation method was success-
fully carried out using a hypothetical speech community con-
taining 15 adults and 5 children. [21] The behavioral
model was extremely simple, as were the grammars (limited

to a tiny subset of English). The learning model was also simplistic, involving the actual borrowing of full-fledged rules rather than their synthesis from fundamental analytic heuristics. The goal of this test, $_{to}$ attain linguistic and social stability through several generations, was attained. It was important because it demonstrated control of the model as a preliminary to innovations that might introduce linguistic or social change. (The particular simulation used a different kind of phrase structure rule notation than we currently use.)

Now our research is directed toward the testing of the methodology through simulation of language change in a real speech community, in sufficient depth and detail that the predictions of the simulation will be subject to emprical verification. In our preliminary search for a suitable test case we first selected the speech community on the island of Tikopia in the South Pacific. This community seemed ideal because of the existence of excellent functional ethnological studies by Raymond Firth that took place in 1928-29, 1951 and 1962 [7-11], and because Tikopia was virtually untouched by World War II. Both the pertinent detail of Firth's studies and the relatively restricted and documented foreign contacts during this period seemed ideal for our work, and we put some effort into designing a simulation system that could handle Tikopian Society and yet

have a basic generality. Unfortunately, Firth was unable
to supply us with his linguistic field notes for Tikopia
(little else of a suitable nature exists).

We then decided to switch to a simulation of language
change among the Maori of New Zealand. The documentation
for this group is voluminous and covers several centuries.
Of particular value is the existence of census data on the
Mario dating back to the nineteenth century. The time
scale and detail level of the Maori model must
be of a coarser sort than for Tikopia because of computer
time and space demands, for it must account for a popula-
tion 40 to 90 times greater than that of Tikopia over a
time period of perhaps 150 years. However, we found that
the design of our simulation system needed little or no
modification for the Maori.

We explicate the representation of both soci-linguistic
situations in Section 3 to provide the reader with insight
into the methodology.

2. Language Learning Component

The language learning logic of the AUTOLING System will
furnish the basis for the learning component of the simula-
tion system. AUTOLING is an automated linguistic field-
worker capable of learning generative grammars through tele-
type interaction with a live human informant. The program

is operational on the Burroughs 5500 computer*, and has

been successfully tested on selected problems in English,

Latin, Roglai, Indonesian, Thai and German. The discovery

methods are heuristic rather than algorithmic, and the

system is under continued modification. One subcomponent

is capable of learning context free phrase structure rules

in response to informant inputs consisting of sentences

segmented into morphemes. An attempt is made to parse each

informant input sentence on the basis of the current tenta-

tive grammar. If the rules are adequate, the program prints

the fact in a teletype message. If not, it posits rules

that might enable the parsing process to be completed.

These rules, and their more general ramifications for the

grammar as a whole, are tested via productions offered to

the informant for acceptability verification. Rejected

sentences cause the newly posited rules to be discarded.

Acceptance of false rules through incomplete testing can

occur. At the present time, the program tests for such a

possibility by attempting to parse various known illegal

sentences. The most recently recorded ones are tested every

time a new rule is coined. All illegal sentences are tested at

periodic intervals. If the bad rules were coined too far in the

---

*Preliminary programming is in ALGOL for the Burroughs
5500 computer, eventually the program will be shifted to
the compatible Burroughs 6500.

past for correction, the program throws out its entire grammar, and reanalyses the entire corpus, using the illegal sentence responsible for the situation as one of the key controls on the new grammar. A later version of AUTOLING will make a stronger attempt to determine the specific culprit rules, and take corrective action in form of transformations or simple context-sensitive phrase structure rules. In fact, eventually, the system will learn a transformational grammar consisting of unordered phrase structure rules plus obligatory transformations that operate whenever conditions permit during the generation process. Also, a morphology learning component will be integrated into the system.

For the simulation system, the human informant is replaced by another grammar associated with another member of the community. While the system will contain only one learning program with its associated parsing and generation routines, each grammar associated with each member of the community might, on various occasions, serve as the grammar in which learning takes place, or as the grammar used to accept or reject the productions of an 'embryonic ' grammar. Learning feedback in an adult-adult conversation will not occur as often as in a child-adult context. The exact circumstances under which an individual's grammar learns or teaches are determined by the socio-demographic model.

Special features that must be added to, or modified in
the AUTOLING system include the following:

a. Multilingual Dictionary: For Tikopia, a list of
Tikopian, English and Melanesian Pidgin morpheme equivalents.
Any individual auditing new lexical items will add a list
link in his grammar (which references terminal elements
only indirectly) to the appropriate entries. Links to
corresponding morphemes (if any exist) in other languages
will be entered only if the person has actually been exposed
to the form in conversation.

For specialized vocabulary, the entries will also con-
tain markers of the context in which the item is to be used.

b. Sentence Generator: Both the Generator and the Parser
use the same grammars. The generator selects non-terminal
rewrite rules according to relative frequency parameters
that are modified during the parsing process. Terminal
elements are referred to by links to the dictionary. Some
terminals are selected on the basis of the generation con-
text, i.e., specialized vocabulary referring to items of
material culture. Under some conditions, a terminal's
translation equivalent in another language may be chosen.

In the specialized case of normative learning, e.g., in
a child-parent relationship, the generator will test newly
formed rules by pertinent test productions offered to the
normative teacher for acceptance or rejection.

c. <u>Parser:</u> The parsing component may modify the fre-
quency parameters pertinent to the generation process as a
function of a particular rule's use in recent parsings.

3. Modelling Tikopia and Maori

Some generality in the system design would be necessary
even if one intended to model only one society. In parti-
cular, the rules governing the interaction of members of
the population would undoubtedly be subject to frequent
revision during the course of research as it might become
apparent that some variables modelled were not pertinent,
and that ommitted ones were significant. A fully general
system, capable of modelling any society, must contain,
implicitly, a universal theory of socio-linguistic behavior.

A basic assumption of our system is that an individual's
group memberships constitute the major determinants of his
conversation behavior. Therefore it is essential that the
system provide an efficient means of describing an individ-
ual's age, sex, political, kin, work and social group member-
ships as well as data of a purely geographic nature.

Specifically, for Tikopia, it seemed that age, sex,
village, clan, religion, household, marital state, work
groups, and social status were the key variables governing
conversational interaction .   We planned to simulate
a thirty or thirty-five year time span in a model contain-
ing a population sample of about 120-165 people distributed

among three villages, representing about 1300 to 1800 people
distributed among approximately 25 villages.

The decision to construct the model with a few villages
containing a large fraction of their real-world population
(as opposed to more villages with fewer modeled people per
village) was made on the basis of material contained in
Firth [7-11] indicating the village as the
largest pertinent unit for our purposes. The decision to
model three villages was based on the recognition of the
subordinate, but real pertinence of inter-village relations.
The problem of representing a complete multi-generation
kinship structure for each individual also set a lower bound
on the number of people per village.

The actual method of creating an initial population
state is rather complex, and is described in Section 4.

The researcher attempting to model Maori society is
faced with the problem of finding pertinent data in a vast
literature of essentially non-pertinent material. Fortu-
nately, official government census information, dating back
to the mid-ninteenth century, provide valuable demographic
data.

The population size demands a different kind of sampling
than in the Tikopia model. The population ranged from
56,000 in 1857-8 to 167,000 in 1961. A study of the litera-
ture suggests that the Maori-English linguistic acculturation

phenomena might best be modelled in the following way:

    a. Population: Sample size ranging from 100 to 300 Maori plus English speakers.

    b. Geographical Distribution: Two communities remote from white contact, plus the graduate creation of a city population group, and a group in an intermediate location.

    c. Key Social Variables: Tribe, hapū or tribal subgroup, social class (aristocrat or commoner), age group (child, young unmarried, young married to middle aged, elder) lineage, work groups or occupation, and religion. The hapū, rather than the immediate family, appears to be the minimal significant social unit of organization for the goals of our simulation. In the case of city dwelling Maori, residence in the same city constitutes another group membership.

    d. Meta-model of Historical Change: Gradually increasing contact with English speakers, wars, gradual migrations to urban areas.

4. System Organization and Construction of the Data Base

The learning program, as it stands, demands an interaction between a live informant and teletype outputted questions. It is necessary, for the purpose of reducing the enormous computer time required for the successful

simulation of change in linguistic patterns, to be able to
break the current program into two parts -- one part that
can read sentences input to it without asking for immediate
help, and another which will generate sentences randomly,
based on the rules that were formulated during the input
stage.

The portion of the program that is responsible for the
generation of random sentences will also determine the con-
text in which the sentence was spoken. Context is deter-
mined by defining the subclass of persons who would be
listening to this sentence, and placing an indicator of this
subclass in the file of sentences which are generated. The
sentences will be placed in a file, that will later be
passed against all individuals in the sample in order that
particular aquaintances are able to "hear" what was said--
at the same time creating rules which shall be used in the
next generation pass.

At major points in the process, events take place that
need not be thought of during the normal cyclic activity.
These involve the life and death routines, marriage cere-
monies, arrival-departures, and recreation of the aquain-
tance lists that describe who is listened to. Because of
the one-to-many character of speeches, it is possible to
keep the aquaintance lists to a manageable size by listing
only those persons whom one listens to, and not those who

are spoken to.

Before we examine the conversation process further let
us discuss the general problem of creating a sample for
data that is available only in aggregate form.[*]

4.1 Sample Generation

For many groups to be studied by the process described
in this paper, samples do not exist. If any information
exists at all about these groups it is often in the form
of cross-tabulation tables published as an indication of
census patterns, and is usually not given in its raw form.

---

[*]The problem of creating a kinship structure is not of
this type. In the case of Tikopia it is essential to keep
track of kin relations with contemporaries that may owe
their origin to links with common ancesters, perhaps 2 or
3 generations removed, who may be deceased at the start of
the simulation. The best automated method we could devise
involves running an accelerated, prefatory, partial simu-
lation of the society beginning several generations before
the official start date. The only aspects modelled would
be those governing birth and death, residence change and
marriage rules. Initially, all individuals would be assumed
to be unrelated, and marriage would take place with rela-
tive freedom. As the prefatory simulation progresses through
successive generations, kin ties are created, and the free
choice of spouses disappears. By the time the presimulation
is completed, the original starting population is dead, and
each member of the main simulation population has a complete
and consistent set of kinship relations. The level of de-
tail in the Maori situation does not demand this micro-
computation of kinship (see Section 3).

To model groups of people where it is impossible to collect
raw data because of expense, time, or other complications
such as the passage of time rendering the sample change
(historical groups), it is often necessary to create a
sample of people artificially. Since any such attempt will
result in an incorrect sample, it is important to realize
this beforehand and be on guard when viewing the results
of the study against arriving at conclusions which are in-
valid. We can, however, obtain results that have some
validity by restricting our discussion to those character-
istics of the sample that we are able to insert into our
sample creation process by the heuristic methods described
below. We realize that heuristic processes are just that--
there is no real guarantee of success in creating a sample
which is totally accurate. But by prefacing the results
of our study with this disclaimer, and restricting our
stated conclusions to those population characteristics
which we know to be true, useful research can be expected.

We can illustrate the sample creation best by an immediate
example.

Suppose we are interested in the study of linguistic
patterns as they are formed with respect to three variables--
age, sex, and marital status. It is necessary first to des-
cribe the catagories that are important to us for
each of the variables in the model.

If we posit that age does not influence linguistic patterns except in major catagories, we can break the ages into the three groups Young, Adult, and Elder.

Since the other two variables Sex and Marital Status have well defined groupings (Male, Female; Married, Unmarried), we can define our task with the following table:

Age	Sex	Marital Status	Percent of Population
Young	Male	Married	?
Adult	Male	Married	?
Elder	Male	Married	?
Young	Female	Married	?
Adult	Female	Married	?
Elder	Female	Married	?
Young	Male	Unmarried	?
Adult	Male	Unmarried	?
Elder	Male	Unmarried	?
Young	Female	Unmarried	?
Adult	Female	Unmarried	?
Elder	Female	Unmarried	?

Defining a population artificially for the requirements of the simulation process involves the accurate choice of percentages of the total population for each of the above permuted catagories of variables. This can be done in many ways.

1. By hand. The above percentages may be chosen by
the researcher after careful reading of documents descri-
bing population characteristics.

2. By computer algorithm. There are often published
statistics on populations that can be used to create appro-
priate percentages. Cross-tabulation tables are the most
fruitful in this attempt, as they often contain all of the
necessary information within them. If they do not, other
population statistics such as correlation matrices may be
used (e.g., lacking a published table displaying the rela-
tionship between Age and Income, a correlation coefficient
of .46 is useful). Since some of the information may be
either contradictory or of disproportionate value, it is
necessary that a decision be made on the actual distribution
characteristics. If tables are available showing the rela-
tionships, they should be used. But if tables are not avail-
able, or if the only available information about a particular
relationship is in the form of another statistic, the prefer-
able thing to do is to create the table by hand, based on
research of the textual material.

For example, assume that we wish to build a file of per-
sons as mentioned earlier. In reviewing the published tables,
however, we cannot find a table relating Age and Marital
Status. We do find, on the other hand, that the correlation
between Age and Marital Status is given as .43. Using this

information, together with research of the text, it may be
possible to generate a table of the following form:

	Married	Unmarried
Young	2%	28%
Adult	46%	16%
Elder	3%	5%

If we make no use of the knowledge of the correlation
coefficient of .43 between Age and Marital Status, we may
generate a sample that has serious faults. Not making
use of it in this case would be similar to creating a table
of the form:

	Married	Unmarried
Young	15%	15%
Adult	32%	30%
Elder	4%	4%

(approximate correlation coefficient $r = 0$)

This table is clearly incorrect. Marital Status should
not be distributed evenly with respect to age.

If a process of random selection over the specified
probability distributions (the relative frequency tables)
is used to create the persons in our sample, it should be

possible to run a cross-tabulation on this data with the
result being that we can reproduce the tables that we
started with to create that data.

After the process of sample creation is finished, we
may produce a table of the form:

Age	Sex	Marital Status	Percent of Population (rounded)
Young	Male	Married	0%
Adult	Male	Married	22%
Elder	Male	Married	2%
Young	Female	Married	2%
Adult	Female	Married	24%
Elder	Female	Married	1%
Young	Male	Unmarried	13%
Adult	Male	Unmarried	9%
Elder	Male	Unmarried	2%
Young	Female	Unmarried	15%
Adult	Female	Unmarried	7%
Elder	Female	Unmarried	3%

This table would imply the use, in part, of the marginal
relative distributions:

```
Male       48%
Female     52%

Young      30%
Adult      62%
Elder       8%

Married    51%
Unmarried  59%
```

It can be seen that since there are few (the number is
rounded to 0%) young married males, more information was
used to arrive at these values than merely the use of the
marginal distributions. Their use alone would imply that
there should be approximately $6\frac{1}{2}$ % young married males.

4.2 Aquaintance Lists.

To model the linguistic patterns as they occur in the
real world, it is necessary to account in some way for appro-
priate dissemination of information by insisting that each
person speak for the most part with the same persons he
spoke to in the past. This is a tedious process if done
dynamically at the time the conversations are to take place
in the computer simulation. We can show that it is par-
simonious to create an "acquaintance list" of those persons
who are in frequent contact with each individual, and to change
this acquaintance list at more infrequent intervals. The
acquaintance lists may be updated together with other major
actions, such as the birth and death routines, arrivals and
departures, and the occurance of natural phenomena such as
seasonal change.

We may build the acquaintance list by a technique close-
ly approximating that which occurs naturally by the "best
fit" method in which two persons are said to be "acquain-
tances" if they have various attributes in common -- they
may live near each other, work together, or belong to the
same social group. If many attributes are in common, then,
these people will be very likely to be forced to speak to
one another whether or not they might be classified correct-
ly as "friends".

More formally, we may define a person's attributes by
his position in the sample space. For a sample of n
variables, a person can be defined by the n-tuple $(V_1, V_2, \ldots,
V_n)$. By a simple calculus, we can map this point from the
integer n-space into the boolean m-space, where m is
greater than or equal to n , and each variable now has the
value 1 if the persons can be characterized by the truth
of this attribute, and 0 otherwise. For example, the
variable Age, in our example above, would be changed from one
variable with three values to three variables with two
values each. From Age: 1=Young, 2=Adult, 3=Elder, we
would construct Young in Age: 1=True, 0=False; Adult in
Age: 1=True, 0=False; Elder in Age: 1=True, 0=False.

A person in our sample can now be characterized by the
boolean m-tuple $(B_1, B_2, \ldots, B_m)$. In order to determine which
attributes that two persons have in common, it is necessary

to ADD (multiply) these to boolean vectors together. The
resultant vector has 1's in the positions where the two
persons origionally both had had 1's, and no place else
are there 1's.

To account for the disproportionate import of the fact
that two attributes are in common, and in some instances
to correct for the fact that persons may be more likely to
be acquaintances if they do not have two particular attri-
butes in common (e.g., Sex), the resultant vector is multi-
plied by a third Weight vector $\overline{W}$.

The Resultant vector is summed to a scalar, and this
number is compared to an externally specified "hit" value
"H" to determine whether these two persons are said to be
"acquaintances". Example:

$$\overline{A} = (1,0,0,1,0,1,1,0,1,0,0,0,0,1,0,1,1,0,0,0,0,1,1)$$

$$\overline{E} = (0,1,1,1,1,0,0,0,0,0,0,0,0,0,1,1,0,0,0,0,1,0,1)$$

$$\overline{AB} = (0,0,0,1,0,0,0,0,0,0,0,0,0,0,0,1,0,0,0,0,0,0,1)$$

$$\overline{W} = (4,4,5,4,5,4,4,4,-5,4,4,4,4,4,4,4,1,4,4,4,4,4,4)$$

$$R=\overline{A}\overline{B}\overline{W} = (0,0,0,4,0,0,0,0,0,0,0,0,0,0,0,4,0,0,0,0,0,0,4)$$

$$X = \sum_{i=1}^{m} R_i = 12$$

$$H = 35$$

In this case we see that since the value of our calcu-
lation X does not exceed or equal the hit value H , we
reject these two persons as being acquaintances. This re-

jection can be easily changed into a more dynamic technique
by the use of more sophisticated stochastic methods, such
as the rejection being conditions on a random number ex-
ceeding the difference between the numbers X and H .
Further selection is necessary to determine one-sided re-
lationships. It may be possible that A is an acquaintance
of B (B listens to A) but B is not an acquaintance of
A (A does not listen to B), for instance if A is a
village chief, and B is a non-destinguished village mem-
ber.

4.3 Conversation Interaction

The flow of the generation and parsing process is as
follows (the only exceptions are in the case of normative
learning were immediate auditor feedback is required):

    a. Conversation Creation:

        1. Generate all utterances from each grammar at
            one time, by passing the grammar file serially.

            A. The number of utterances for each pass is
               set as an external parameter "S" .

            B. Generate "S" sentences in context C_1 .
               Generate "S" sentences in context C_2 .
               Generate "S" sentences in context C_3 .
               ⋮
               Generate "S" sentences in context C_n .

        2. Enter conversation creation routine.

The conversation creation routine will peruse the acquaintance lists of each persons to generate "listens" in the form of ordered triples (a,b,c), where

    a = ID of listener

    b = ID of Acquaintance

    c = Context of talk.

This triple (a,b,c) will be placed in a file called the "listen" file. The "listen" file, when finished, will be a stack of entries in order by the first entry a_i .

$$a_1 \ b_1 \ c_1$$
$$a_2 \ b_2 \ c_2$$
$$a_3 \ b_3 \ c_3$$
$$\vdots \ \ \vdots \ \ \vdots$$
$$a_n \ b_n \ c_n$$

3. Enter the routine which parses the sentence produced.

b. The Parsing Process:

For any two persons A and B, A can listen to sentences produced by B in only one context.

   1. Bring in the grammar for person A from second-level memory

   2. Determine the address on second-level memory of the conversation specified by the triple (a,b,c)

and bring it into first-level memory (core).

3. Parse, or "listen" to the sentence.

4. Iterate on step 2 until all sentences are parsed.

5. Put the new grammar for this person on second-level memory.

6. Get the next grammar from second-level memory and go to step 2.

7. If no next grammar, increment the time counter.

8. If time to recreate the acquaintance lists or other major events such as birth/death routines and arrivals/departures do so.

9. Iterate on step 1 until finished with entire simulation process.

5. Interpretation of Results

The key problem is determining the success or failure of a simulation. Assuming everything else has gone well, how does one compare the grammars of the population members to determine their mutual similarities and their relation to the language situation in contemporary, real world Maori Society?

The design of the system offers a unique, detailed, quantitative method for determining the similarity of the competence of speakers. Every legal sentence ever generated in the course of the simulation is saved by the system. At

the end of the simulation (or some other time) each individ-
ual must attempt to parse every legal sentence ever produced.
Different individuals may expect to have varying degrees of
success in their parsing attempts. Analysis of the results
can offer a detailed, objective picture of the dialect situ-
ation on the basis of common success or failure in parsing
particular sentences. These results may be correlated with
any socio-demographic factors recorded in the data base of
the model.

Given these rssults, one may then send the same list of
sentences to New Zealand, and have the analogous test per-
formed on a sample of the Maori population, asking informants
to indicate the legal and illegal sentences.

The results of the live testing may then be compared
with the simulation results. Thus, the Monte Carlo simu-
lation approach appears to offer Linguistics a strong
empirical methodology for testing otherwise unverifiable
hypotheses.

References and Bibliography

1. Ausubel, David: Maori Youth, Price Milburn, Wellington, 1961.

2. Best, Elsdon: The Maori, 2 vols., Memoirs of the Poly- nesian Society, No. 5, Wellington, 1924.

3. Beaglehole, Ernest and Pearl: Some Modern Maoris, New Zealand Council for Educational Research, Wellington, 1946.

4. Biggs, Bruce: Maori Marriage, Polynesian Society Maori Monographs No. 1, Wellington, 1960.

5. Biggs, Bruce: The Structure of New Zealand Maori, Re- printed from Anthropological Linguistics, Vol. 3, No. 3; English-Maori Dictionary, Reed, Wellington, 1966.

6. Buck, Peter: The Coming of the Maori, Maori Purposes Fund Board and Whitecombe and Tombs Ltd., Wellington, First Edition 1949.

7. Firth, Raymond: Economics of the New Zealand Maori, Government Printer, Wellington, First Edition 1929, Second Edition 1959.

8. Firth, Raymond: We the Tikopia, George Allen & Unwin, London, 1936.

9. Firth, Raymond: 'A Note on Descent-Groups in Polynesia', Man, Vol. 57, 1957, pp. 4-8.

10. Firth, Raymond: Social Change in Tikopia, George Allen & Unwin, London, 1959.

11. Firth, Raymond: Primitive Polynesian Economy, 2nd Edition, Routledge & Kegan Paul, London, 1965.

12. Groube, L. M.: Settlement Pattern in Prehistoric New Zealand, M.A. Thesis, The University of Auckland, 1964.

13. Hargreaves, R. P.: 'Changing Maori Agriculture in Pre- Waitange New Zealand', J.P.S*, Vol. 72, No. 2, 1963, pp. 101-17; 'Maori Agriculture of the Auckland Province in the Mid-Nineteenth Century', J.P.S., Vol. 68; No. 2 1959, pp. 61-79; 'Maori Agriculture after the Wars (1871-1886)', J.P.S., Vol. 69, No. 4, 1960, pp. 354-67.

14. Harre, John: 'The Relevance of Ancestry as a Factor in Social and Cultural Choice', J.P.S., Vol. 74, No. 1, 1965, pp. 3-20.

*Journal of the Polynesian Society.

15. Harre, John: Maori and Pakeha, Institute of Race
    Relations and Oxford University Press, 1966.

16. Hawthorn, H. B.: The Maori: A Study in Acculturation,
    American Anthropology Association Memoir 46, 1944.

17. Holst, Halvor: 'The Maori Schools in Maori Education',
    Education, Vol. 7, Department of Education, Wellington,
    March 1958, pp. 53-59.

18. Hunn, J. K.: Report on Department of Maori Affiars
    24 August 1960, Government Printer, Wellington, 1961.

19. Keesing, Felix: The Changing Maori, The Board of Maori
    Ethnological Research, New Plymouth, 1928.

20. Klein, S.: Some Components of a Program for Dynamic
    Modelling of Historical Change in Language. Paper No.
    14 of Preprints of Invited Papers for 1965 International
    Conference on Computational Linguistics, May 19-21,
    New York. [To appear in Russian in a book on Language
    Contact edited by V. Rosentsveig].

21. Klein, S.: Historical Change in Language Using Monte
    Carlo Techniques. Mechanical Translation, Vol. 9,
    Nos. 3 & 4, September & December, 1966 (Pub. May 1967).
    [To appear in English in an English version of the
    above mentioned Rosentsveig book.]

22. Klein, S.: Current Research in the Computer Simulation
    of Historical Change in Language. In press, Proceedings
    of the Xth International Congress of Linguists, 1968,
    Romanian Government, Bucharest. (See papers presented).

23. Klein, S., Fabens, W., Herriot, R., Katke, W., & Towster,
    A.: The AUTOLING System. University of Wisconsin Com-
    puter Sciences Department. Technical Report No. 43, Sept.
    1968.

24. Maunsell, R.: Grammar of the New Zealand Language,
    Upton & Co. Auckland, 1894.

25. Maori Affairs, Department of: The Maori Today, (a) First
    Edition 1949; (b) Third Edition 1964; Government Printer,
    Wellington.

26. Metge, Joan: 'Marriage in Modern Maori Society', Man,
    Vol. 57, 1957, pp. 166-70.

27. Metge, Joan: A New Maori Migration: Rural and Urban
    Relations in Northern New Zealand, London School of
    Economics Monographs in Social Anthropology No. 27,
    Athlone Press and Melbourne University Press, 1964.

28. Metge, Joan: The Maoris of New Zealand, Rutledge and
    Kegan Paul, London, 1967.

29. Miller, John: Early Victorian New Zealand, Oxford
    University Press, 1958.

30. Parr, C. J.: 'A Missionary Library, Printed Attempts
    to instruct the Maori', J.P.S., Vol. 70, No. 4, 1961,
    pp. 429-50; 'Maori Literacy 1843-1887, J.P.S., Vol. 72,
    No. 3, 1961, pp. 211-34.

31. Sinclair, Keith: The Origin of the Maori Wars, New
    Zealand University Press, Wellington, 1957.

32. Sorrenson, M. P. K.: 'Land Purchase Methods and their
    Effect on Maori Population 1865-1901', J.P.S., Vol. 65,
    No. 3, 1956, pp. 183-99; 'The Maori King Movement 1858-
    1885', Studies of a Small Democracy (edited by Robert
    Chapman and Keith Sinclair), Paul's Book Arcade for
    the University of Auckland, 1963.

33. Williams, H. W.: A Dictionary of the Maori Language,
    6th edition, Revised and Augmented under the Auspices
    of the Polynesian Society, R. E. Owen, Government Printer,
    Wellington, 1957.

34. Williams W. L. & Williams H. W., First Lessons in Maori,
    Trubner, London, 1862.

35. Wright, Harrison M.: New Zealand 1769-1840: The Early
    Years of Western Contact, Harvard University Press, 1959.