Sep 1957 J. Cocke

# AN EXPERIMENT IN MUSICAL COMPOSITION

BY

F. P. BROOKS, JR., A. L. HOPKINS, JR., P. G. NEUMANN, AND W. V. WRIGHT

*Reprinted from* IRE TRANSACTIONS
ON *ELECTRONIC COMPUTERS*
Volume EC-6, Number 3, September, 1957

PRINTED IN THE U.S.A.

For John Cocke
with compliments
Frederick P. Brooks Jr.

*Reprinted from* IRE TRANSACTIONS
ON *ELECTRONIC COMPUTERS*
Volume EC-6, Number 3, September, 1957

PRINTED IN THE U.S.A.

# An Experiment in Musical Composition*

F. P. BROOKS, JR.†, A. L. HOPKINS, JR.‡, P. G. NEUMANN‡, AND W. V. WRIGHT‡

***Summary*—The high-order probabilities of element sequences can be determined from a sample of linear structures and can be used for synthesis of new structures. From theoretical considerations one can identify the qualitative conditions for satisfactory output. The theoretical concepts can be tested and quantitative parameters determined by experiment. Such an experiment has been performed by analyzing written music and by testing the analysis through the synthesis of new musical compositions, using a digital computer.**

**A sample of 37 melodies was analyzed for the probabilities of the elements, element pairs (digrams), trigrams, and so on to the eighth order. The tables derived were used for the synthesis of original melodies by a random process. The theory and the experimental verification are considered in detail. The experimental results presented include comparative statistics of the successful syntheses using each of the eight orders of analysis, examples of melodies generated by low, medium, and high-order synthesis, and confirmation of degeneracy and other effects predicted by the theory.**

## INTRODUCTION

HUMAN thought processes have at least two radically different components, often identified as induction and deduction. Digital computers are readily programmed to perform deductive "reasoning," but their ability to draw generalizations from special cases is extremely limited. Many interesting computer experiments in game playing, "learning," theorem proving, etc., have been aimed at discovering methods of simulating rudimentary inductive processes with a computing machine.[1,2]

---

* Manuscript received by the PGEC, April 20, 1957.

† Product Dev. Lab., International Business Machines Corp., Poughkeepsie, N. Y. Formerly with The Computation Lab., Harvard Univ., Cambridge, Mass.

‡ The Computation Lab., Harvard Univ., Cambridge, Mass.

[1] C. E. Shannon, "Computers and automata," PROC. IRE, vol. 41, pp. 1234–1241; October, 1953.

[2] A. G. Oettinger, "Programming a digital computer to learn," *Phil. Mag.*, vol. 7, pp. 1243–1263; December, 1952.

There are often demands for inductive reasoning where the results of the generalization process do not need to be stated explicitly as rules but only need be in a form suitable for subsequent deductive reasoning. This is the case whenever one attempts to synthesize structures of a certain class, as in the creation of synthetic linguistic utterances or synthetic musical compositions. The simplest way to perform such tasks is for a human to analyze some sample of the type of structure desired, draw up some explicit rules and constraints, and allow the machine to operate deductively, although perhaps at random, in the synthesis of new structures. Some workers have performed computer-implemented musical composition in this manner.

It is of considerably more interest to attempt to synthesize musical compositions by having the machine inductively analyze a sample of acceptable compositions and, using its conclusions, deductively synthesize new but original compositions. Such an induction can be performed by determining the probabilities of note sequences. Several theoretical aspects of such a process deserve examination.

## THEORY OF ANALYSIS-SYNTHESIS

The derivation of sufficient information from a sample to permit subsequent deductive synthesis of new and original members of the same class of structure depends not only upon the sophistication of the analysis, but also upon the characteristics of the sample analyzed. Suppose one undertakes to analyze a small number of tunes and to use the results in a random process for synthesizing new tunes. One can anticipate three causes of difficulty which can be visualized with the aid of the diagram in Fig. 1. This diagram shows that the process of analysis of several members of a sample $S_1$, $S_2$, $\cdots$, $S_n$, belonging to some common class of structure, yields a generalization $G$. From the generalization, a synthesis is used to derive one or more new structures $S'$, which one hopes will be of the same class of structures as the original sample. The first difficulty is that overly naive analysis may yield a generalization so loose that the resulting $S'$ does not belong to the same class of structure as the sample members. For example, synthetic note sequences may not belong to the class of acceptable melodies from which they were derived. The second difficulty is that the sample may be so small that no generalization may be drawn no matter how sophisticatedly an analysis is performed. The third difficulty is that the sample members may be so alike that when the generalization is formed, it is impossible to create a new structure of the same class which is not identical to one of the structures of the sample.

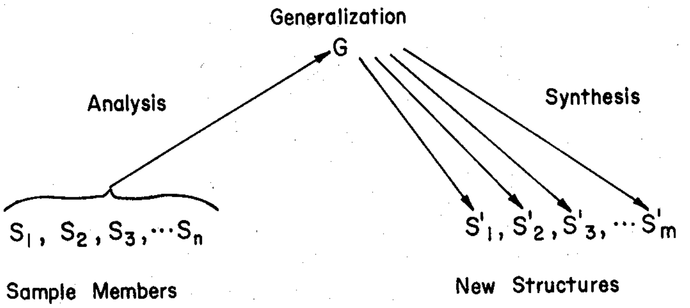

Fig. 1.

It appears worthwhile to perform some experiments to verify the existence of these three effects and to learn something about the basic nature of the analysis and synthesis process. For this purpose, two structure classes of interest are the class of linguistic utterances acceptable to human beings as "meaningful" and the class of musical compositions acceptable to human beings as something a human might have created. The synthetic linguistic utterance problem (which includes the well-known sonnet-writing problem) is more general, more interesting, and more difficult than the musical composition problem. The musical composition problem, interesting in its own right, also serves as a small-scale model of the linguistic utterance problem. For these reasons the experiments described herein were undertaken on the application of the analysis-synthesis process to the musical composition model.

The analysis of any structural system consists of the determination of a set of basic elements and the determination of the combinational relationships among the elements. For the present, we consider that only the second step is sufficiently difficult mechanically to be worth computer mechanization and sufficiently simple conceptually to permit computer mechanization.

The constraints of combination of an element set can be stated in several different ways. The first of these is the *explicit* statement of rules as to what combinations may or may not occur. For example, the explicit method of representing linguistic constraints is illustrated by the schoolboy's grammar and spelling books. Here one finds several specific examples of a connective pattern grouped together with the generalization from these to a rule: "*i* before *e*, except after *c*, or when sounded as *a*, as in *neighbor* or *weigh*." From the set of rules on spelling, the explicit method extends to the formulation of rules for the combination of larger elements in formal grammar, syntax, and rhetoric. The expression of all the combinational constraints in a structure by such a method may grow quite complex with many levels of rules, sub-rules, exception systems, and specific exceptions. Examination of any complete grammar reveals this complexity at which the "*i* before *e*" example only hints.

A second approach is the *exhaustive* one in which all the combinational constraints are listed explicitly. A class can always be completely described by listing its members, and this description is often more useful, more revealing, and less confusing than any set of class characteristics. We do, in fact, learn the constraints of combination for *i* and *e* by long and painful use of spelling book and dictionary which list all the extant combinations. By this means, our description of constraints extends to *seize* and *financier* without any special treatment. While the exhaustive approach has the advantage of conceptual simplicity and guarantees a complete description of any structure, it is tedious and burdensome to apply to any system as complex as language or music.

Both the explicit and the exhaustive methods of constraint description are incomplete; they only indicate which combinations may occur without describing the relative frequency of occurrence. The mere statement that *t* and *e* can each follow an initial *p* is incomplete without the information that *pt-* is much less common than *pe-*. This leads to a third method of constraint description, the probabilistic statement.

The probabilistic description shares the advantages of the exhaustive description; it is simple in concept and uniform in application. Ease and certainty of determining combinational constraints in analysis and following them in synthesis are independent of the complexity or obscurity of the constraints themselves. For computer application these are cardinal virtues, and the great magnitude of the task of formulating and applying such a description is a less important problem. The probabilistic method has a further advantage over other methods: the precision and validity of the description are little affected by the inclusion or exclusion of sample members beyond a sample of a certain size. With both the explicit and the exhaustive methods, the existence of an exception not included in the sample analyzed would compromise the accuracy of the analysis. Finally, although analyses for some purposes are required to yield an explicit statement of the common properties of a sample, this is not required for the analysis-synthesis process shown in Fig. 1.

For these reasons, the probabilistic method was chosen for the analysis-synthesis experiments described below. Analysis by probabilities of combinations may be taken to any order. One may use simple probabilities of occurrence, simple transitional probabilities in a sequential, linear structure, or conditional probabilities depending upon any number of preceding or adjacent elements.

The concept of using probabilistic analysis as part of an analysis-synthesis procedure is by no means new. Shannon[3] proposed a simple and elegant method of performing an implicit analysis and synthesis of linguistic utterances. Pinkerton[4] and F. and C. Altneave[5] have used simple transitional probabilities for the analysis and synthesis of melodies.

The use of probabilistic descriptions of several orders, however, permits one to apply successively more sophisticated analytic procedures to a given sample in an attempt to determine the existence and behavior of the three hindering effects set forth earlier. This was the primary purpose of the experiments described below. It was an incidental but fervent hope that the sample selected would be sufficiently large and rich to permit some of the more sophisticated analysis-synthesis procedures to yield musical compositions that were both original and musically acceptable.

For structures that are linear or one-dimensional, higher-order transitional probabilities are known as *Markoff chains*, and we shall refer to their determination as a *Markoff analysis* and their use in the construction of new structures as a *Markoff synthesis*. Let us examine in more detail the analytic and synthetic methods used.

## METHOD

A Markoff analysis of order $m$ is the process of determining certain joint probabilities of the occurrences of elements in a sample. Thus, within the sample, a given sequence of $(m-1)$ elements may be followed by any one of several elements. We wish to find the relative frequency of each of these elements after the given $(m-1)$ element sequence. A sequence of $m$ elements will be called an *m gram*. A sample will ordinarily be made up of several independent sequences of elements called *utterances*. A step in a Markoff synthesis of order $m$ is performed by choosing an element to follow a given element sequence of length $(m-1)$ in such a way that, if this process were to be repeated a large number of times, the relative frequencies of the chosen elements would match those found in the sample. Since it is desired that the synthesis be random within these restrictions, the choice is made according to the magnitude of a random or pseudo-random number. It is sufficient for this purpose, then, to construct for each element sequence of length $(m-1)$ a table of cumulative probabilities which allots to each $m$th element a segment of the range of numbers between zero and one, the size of the segment being the desired relative frequency of the note. Then a random number between zero and one falls into the range allotted to an element with a probability equal to the relative frequency of the element. This procedure is, in effect, a mapping of a uniform distribution of random numbers into the nonuniform distribution of elements following a given sequence of $(m-1)$ elements. Consider, for the sake of illustration, the following hypothetical fourth-order analysis and synthesis of musical text. Suppose that the occurrences of a three-note sequence $C-E-G$ in the sample are as shown below:

$C-E-G-A$
$C-E-G-A$
$C-E-G-C$
$C-E-G-E$
$C-E-G-E$
$C-E-G-G$
$C-E-G-G$
$C-E-G-G$
$C-E-G-G$
$C-E-G-G$.

---

[3] C. E. Shannon and W. Weaver, "The Mathematical Theory of Communication," Univ. of Illinois Press, Urbana, Ill., pp. 11–15; 1949.

[4] R. C. Pinkerton, "Information theory and melody," *Sci. Amer.*, vol. 194, p. 77; February, 1956.

[5] F. and C. Altneave, unpublished study described by H. Quastler, "London Symposium on Information Theory," Butterworth, Ltd., London, England, pp. 168–169; September, 1955.

The relative frequencies of notes following the sequence $C-E-G$ are:

$A \quad 2/10 = 0.2$

$C \quad 1/10 = 0.1$

$E \quad 2/10 = 0.2$

$G \quad 5/10 = 0.5.$

Synthesis may then be carried out by means of the table of cumulative probabilities which, for any argument, $x$, between zero and one, will determine the note to be selected. One of the several possible arrangements of such a table is:

$A \quad 0 \leq x < 0.2$

$C \quad 0.2 \leq x < 0.3$

$E \quad 0.3 \leq x < 0.5$

$G \quad 0.5 \leq x < 1.0.$

If the random numbers $x$ are uniformly distributed between zero and one, the probability of choosing a $C$ is one tenth; and in the course of synthesis, the sequence $C-E-G$ will be followed by $C$ with expectation of occurrence of one tenth. Clearly, the conditional probabilities found by Markoff analysis can be approximated to any desired accuracy by this method.

The three hindering effects considered above can be understood more precisely in terms of the Markoff analysis and synthesis method. The first effect, overly naive analysis, will show itself when only low-order transition probabilities are used, so that the structure of the synthetic composition is insufficiently constrained. This permits the occurrence of note sequences that are not acceptable as legitimate musical structures.

The second effect, insufficient sample size, shows itself at some order $m$ of analysis and synthesis where each distinct $m$ gram occurs only once. From that order upward, further analysis is useless since the higher-order transition probabilities are completely determined by the analysis to that point. There is no more information to be extracted from the sample. Therefore, if the synthesis at this order fails to yield acceptable and original utterances, nothing can be done, for the sample size is so small that further sophistication of the analysis yields no benefits.

The third effect, insufficient sample diversity, will always show itself before or at the same time as the second effect, and so is of much greater practical importance. Suppose the digrams occurring in the sample have been found, and a trigram analysis is performed yielding one distinct trigram for each distinct digram. Even though there may be many occurrences of each digram and thus of its associated trigrams, an utterance is completely specified by giving the initial digram, for it will specify a trigram whose final two elements uniquely specify another, etc. Further sophistication of the analysis is fruitless beyond this point; all the information in the sample has been extracted. If the analysis to this order is insufficient to yield acceptable and original utterances, nothing can be done; the sample is too redundant to yield more information.

From the theoretical examination of the analysis-synthesis problem, one can see that experiments using Markoff analysis and synthesis will yield any of three results. If too elementary or low-order analysis is used, the results will not resemble the sample members closely enough to be recognized as members of the same class. If too high an order of analysis is attempted for a given sample size and diversity, the synthesized results will degenerate; that is, they will duplicate sample members. Or, if the sample is sufficiently large and diverse, there will be some orders of analysis for which the results are original and still recognizable as members of the class.

For the sample of music selected, each of the three results was found for some order of the Markoff analysis-synthesis.

## THE EXPERIMENTS

As indicated above, both linguistic utterances and musical compositions are interesting structure classes upon which to perform experiments in Markoff analysis and synthesis; the musical structures were chosen because of the greater simplicity. The selection of music did introduce problems, however. Music has many dimensions, such as pitch, meter, rhythm, key, harmony, dynamics, and quality. Some selection must be made as to which to include as experimental variables and which to ignore. Since quality is not expressed in written music, and difference in key can be grossly compensated by transposition, these were not treated as variables. Harmony introduces a vertical as well as a horizontal structure to the text, and the need for simplicity dictated that the present experiments be confined to the horizontal structure of a single line. The metrical structure was fixed as an experimental variable by using a sample all of whose members had a fixed metrical structure.

After written music was chosen as the structure class to be used in the experiments, a sample was assembled. It consisted of 37 common meter hymn tunes, a choice determined largely by the decision to work within one metrical structure. The common meter hymn is perhaps the most widely used rigid metrical structure and one of the simplest. In fact, several variants are found even within this structure, and the sample was confined to the most common single variant. The common meter hymn tune has the additional advantage of providing a fairly large collection of compositions from different composers and centuries of origin. The hymns in the sample all begin on the last beat of a four-beat measure, and none have any notes shorter than an eighth note.

Several analyses of the sample and syntheses of new hymn melodies were performed with a large-scale computer. The machine has random-access storage of 230 sixteen-decimal digit numbers and separate magnetic drum storage of 4000 numbers and 10,000 instructions. It is synchronous in operation with most instructions requiring 1.3 μsec.

For computer manipulation it was necessary to encode the notes as numbers. A range of four octaves in the chromatic scale was selected, and each tone was assigned a two-digit number. Different time values were represented by dividing the whole hymn into 64 eighth-note cells. The content of each cell was either a tone struck in that cell or one held over from the preceding cell. The even integers from 02 to 98 were used to represent a struck tone while the corresponding odd integers from 03 to 99 were used to represent a continued tone. Henceforth, each of the eighth-note cells will be referred to as a note regardless of whether it represents a struck or a held note. In order to provide a common basis for the analysis, each hymn was transposed upwards to the nearest key of *C*.

As each storage location was capable of holding sixteen decimal digits and hence eight notes, it was convenient to carry out analyses up to the eighth order. As it turned out, this order was the one at which redundancy-caused degeneracy became noticeable.

The first step in the analysis was to isolate all of the eight-note sequences or octograms occurring in the sample, and then sort them into numerical order.[6]

This sort placed all of the octograms in order within their initial heptagram, all heptagrams in order within their initial hexagrams, and so on. Hence, the sorted sequences of lower orders were readily obtainable from the sorted octograms by shifting all the octograms the appropriate number of places to the right, thus eliminating the extraneous notes.

To complete the analysis, cumulative relative frequency tables of the type previously discussed were formed for each value of *m* for all of the $(m-1)$-note sequences present in the sample. The relative frequency for each distinct *m*-gram was calculated directly from the sorted *m*-grams by counting the number of occurrences of the *m*-gram. This number; *i.e.*, the number of times the *m*th note followed the initial sequence of $(m-1)$ notes, was divided by the total number of occurrences of all *m*-grams having the same initial $(m-1)$-note sequence, thus giving the relative frequency of the *m*th note with respect to that initial sequence. The entries in the table were calculated by accumulating the relative frequencies pertaining to the same $(m-1)$-note sequence.

As an illustration of this procedure, consider the hypothetical sequence of sorted octograms shown in Table I. The resulting octogram table entries are the distinct octograms (denoted with an asterisk) and their associated cumulative probabilities, shown in the last column. The octogram 3637 2627 3233 2622 corresponds to the eight-note sequence $G\overline{G}C\overline{C}A\overline{A}CD$, where a bar over a letter indicates that the note is held.

TABLE I

SORTED OCTOGRAMS, ILLUSTRATING THE FORMATION OF THE CUMULATIVE PROBABILITY TABLES

| Cell | | | | | | | | Octo-gram Count | Hepta-gram Count | Relative Fre-quency | Cumula-tive Probability |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | | | | |
| 36 | 37 | 26 | 27 | 32 | 33 | 26 | 22 | 1 | | | |
| *36 | 37 | 26 | 27 | 32 | 33 | 26 | 22 | 2 | | 2/8 | 2/8 |
| 36 | 37 | 26 | 27 | 32 | 33 | 26 | 27 | 1 | | | |
| 36 | 37 | 26 | 27 | 32 | 33 | 26 | 27 | 2 | | | |
| 36 | 37 | 26 | 27 | 32 | 33 | 26 | 27 | 3 | | | |
| 36 | 37 | 26 | 27 | 32 | 33 | 26 | 27 | 4 | | | |
| *36 | 37 | 26 | 27 | 32 | 33 | 26 | 27 | 5 | | 5/8 | 7/8 |
| *36 | 37 | 26 | 27 | 32 | 33 | 26 | 28 | 1 | 8 | 1/8 | 8/8 |
| 36 | 37 | 26 | 27 | 32 | 33 | 32 | 33 | 1 | | | |
| *36 | 37 | 26 | 27 | 32 | 33 | 32 | 33 | 2 | 2 | 2/2 | 2/2 |
| *36 | 37 | 32 | 33 | 32 | 33 | 32 | 33 | 1 | 1 | 1/1 | 1/1 |
| 36 | 37 | 32 | 33 | 32 | 33 | 36 | 37 | 1 | | | |
| *36 | 37 | 32 | 33 | 32 | 33 | 36 | 37 | 2 | | 2/3 | 2/3 |
| *36 | 37 | 32 | 33 | 32 | 33 | 36 | 42 | 1 | 3 | 1/3 | 3/3 |

For each *m* the number of these tables was equal to the number of distinct $(m-1)$-grams. Table II, next page, gives the total number of *m*-grams beginning with struck notes, held notes, and initial rests (00). With 64 octograms in each hymn, the 37 hymns yielded 2368 octograms. As seen from the table, only 1701 of these were distinct. Since the number of distinct heptagrams was 1531, there were exactly that many octogram tables having 1701 entries in all. The 4000-word drum storage of the computer thus sufficed for the storage of both the sorted distinct octograms and their associated cumulative probabilities.

## THE SYNTHESIS

Syntheses of orders one through eight were then accomplished by the process of Markoff synthesis discussed earlier, using eight digit pseudo-random numbers obtained from an algorithm known to be nonrepetitive for the first $(10^8+1)/17$ numbers.[7] The first note of the hymn was found by entering the probability table for the *m*-grams whose first $(m-1)$ notes were rests and choosing at random among these. The second note was selected by choosing an *m*-gram beginning with $(m-2)$ rests followed by the first chosen note. The procedure of choosing one *m*-gram, given the preceding $(m-1)$ notes, was carried on in similar fashion until the 64 notes comprising one hymn had been generated. The encoding of initial rests preceding each hymn thus per-

---

[6] The sort was performed by a general purpose digital sorting routine that counts digital occurrences in the $(n+1)$st column while sorting on the *n*th column. Rapid and efficient in use of storage, the routine sorts up to 2000 numbers without requiring any intermediate input or output. The authors are indebted to A. S. Goble III and J. Hines for the programming of the routine.

[7] D. H. Lehmer, "Mathematical Methods in Large-Scale Computing Units," *Proceedings of a Second Symposium on Large-Scale Digital Calculating Machinery (1949)*, Harvard Univ. Press, Cambridge, Mass.; 1951.

TABLE II

DISTINCT *M*-GRAM COUNTS, BEGINNING WITH SEVERAL DIFFERENT NOTES

| Initial Note | Order of Analysis *m* | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 12 $G$ | 1 | 3 | 15 | 25 | 57 - | 66 | 99 - | 102 |
| 13 $\overline{G}$ | 1 | 12 | 30 | 67 - | 84 | 124 - | 131 | 150 |
| 16 $F$ | 1 | 6 | 14 | 25 | 43 - | 45 | 63 - | 63 |
| 17 $\overline{F}$ | 1 | 8 | 14 | 32 - | 33 | 51 - | 51 | 66 |
| 18 $E$ | 1 | 6 | 19 | 32 | 74 - | 83 | 131 - | 131 |
| 19 $\overline{E}$ | 1 | 11 | 23 | 68 - | 78 | 130 - | 135 | 156 |
| 22 $D$ | 1 | 4 | 19 | 30 | 66 - | 69 | 104 - | 109 |
| 23 $\overline{D}$ | 1 | 12 | 21 | 57 - | 64 | 102 - | 106 | 126 |
| 26 $C$ | 1 | 6 | 18 | 28 | 65 - | 71 | 111 - | 112 |
| 27 $\overline{C}$ | 1 | 13 | 25 | 65 - | 72 | 112 - | 113 | 136 |
| All Struck Notes | 18 | 47 | 152 | 219 | 444 - | 479 | 698 - | 705 |
| All Held Notes | 18 | 110 | 182 | 428 | -485 | 717 - | 738 | 869 |
| 00 Initial Rest | 1 | 5 | 10 | 28 | 45 | 70 | 95 | 127 |
| Total Distinct *m*-Grams | 37 | 162 | 344 | 675 | 974 | 1266 | 1531 | 1701 |

TABLE III

PERCENTAGE OF ATTEMPTS YIELDING ACCEPTABLE HYMNS

| Metric Constraint | Order of Analysis *m* | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| Even Quarters | 100 | 40 | 32 | 10.8 | 9.2 | 3 | 2.5 | 8.4 |
| Quarters with Various Options | | 41 | 26 | 15 | 11 | 2 | | 10 |
| Dotted | 100 | 13.5 | 5 | 2 | 1 | 0 | 0 | 0 |
| Dotted with Various Options | 100 | 17 | 4 | 1.5 | 0 | 0 | 0 | 0 |
| Skeleton | 100 | | | 5 | 4 | 3 | 2.4 | 10 |

TABLE IV

EXAMPLE OF *M*-GRAM EXTENSION PROPERTIES

| *m* | Note Brought in | *m*-Gram, Begun with Struck Note | *m*-Gram, Begun with Held Note |
|---|---|---|---|
| 4 | $E$ | $C\ \overline{C}\ D\ \overline{D}$ | $\overline{C}\ D\ \overline{D}\ E$ |
| 5 | $E$ | $C\ \overline{C}\ D\ \overline{D}\ E$ | $\overline{C}\ D\ \overline{D}\ E\ \overline{E}$ |
| 6 | $F$ | $C\ \overline{C}\ D\ \overline{D}\ E\ \overline{E}$ | $\overline{C}\ D\ \overline{D}\ E\ \overline{E}\ F$ |

mitted the synthesis process to operate uniformly, even in starting.

The selection of notes was subject to certain externally applied constraints in addition to those implicit in the frequency tables of the analysis. These explicit constraints, however, were just those permitted by the selection of a sample of uniform metrical structures. In order to force the synthesis to stay within the selected metrical structure, a metric constraint was applied to certain critical notes. For example, the first note of each measure had to be struck rather than held. In addition, the two main phrases both had to end on a dotted half note; hence notes 28 through 32 and 60 through 64 were constrained to be held notes. Each note generated was examined and compared with the metric constraint, if any, for its position in the time basis. If the constraint was not met, the note was rejected and another note was generated by the next random number. If the constraint could not be met after 15 trials, the hymn was discarded and a new one begun. Using these metrical constraints in this manner can be shown to be equivalent to determining and using separate note combination constraint tables for each time point or for subsets of the time points. The rigid metrical structure represented by the sample permitted the separation of metrical constraints from others with a considerable resultant simplification.

An additional constraint was added in the last meassure of the hymn. Since after transposition all of the hymns in the sample ended in *C* above middle *C*, every generated hymn was required to end the same way. If, after fifteen trials, no final *m*-gram was chosen which led into a dotted half note *C* above middle *C*, the first 57 notes already generated were used to begin another synthesis. This device distorted the absolute size of the yield percentages (number of acceptable hymns completed over the number of starts). It did not, however, affect the faithfulness with which synthesized hymns obeyed all combination probabilities, nor should it have affected the relative sizes of the yield percentages as metric constraint or analysis order changed. Its use permitted the production of a significant number of acceptable tunes within a reasonable time.

## RESULTS AND CONCLUSIONS

The synthesis just discussed was carried out with various metric constraints for all values of *m* from one to eight. The constraints were of three types. One of these types was based on even quarter notes throughout the hymn, except during the phrase ends. A second was based on a dotted rhythm in the second and sixth measures. Both of these constraints were varied by the introduction of optional rhythms. The third type was a skeletal constraint requiring only that the first and seventh cells in each measure contain struck notes and that the two main phrases end as usual, leaving the remaining rhythmic structure optional.

In all, over 600 complete hymns were synthesized in some 6000 starts. In cases where the yield was a very low percentage of the number started, more attempts were made than in cases where the yield was high. A summary of the results is shown in Table III, which gives the yield for each order *m* under the various constraints.

The yield is the percentage of the completed hymns out of the total number started. In particular, the similarities between the yields for orders four and five and between the yields for orders six and seven are worthy of comment since they resemble those between the distinct *m* gram counts in Table II. As predicted from theoretical considerations, the transition from an odd *m* to the next higher (even) *m* introduces more implicit constraints than does going from an even *m* to the next higher (odd) *m* because of the basic quarter-note structure. Thus,

TABLE V

METRIC CONSTRAINTS OF EXAMPLES 1–5

|  |  |  |  |  |  |
|---|---|---|---|---|---|
| Ex. 1 | X — | X — O X X — X X | X — X — X — X — | X X X X X — X X | X — — — — — |
| Ex. 2 | X — | X — — X X — X — | X — X — X — X — | X — O X X — X — | X — — — — — |
| Ex. 3 | X — | X — X — X — X X | X — X — X — X — | X — X — X — X — | X — — — — — |
| Ex. 4 | X — | X — X — X — X — | X — X — X — X — | X — X — X — X — | X — — — — — |
| Ex. 5 | X O | X O O O O O X O | X O O O O O X O | X O O O O O X O | X — — — — — |
| *t*; *t*-32 | 1 2 | 3 4 5 6 7 8 9 10 | 11 12 13 14 15 16 17 18 | 19 20 21 22 23 34 25 26 | 27 28 29 30 31 32 |
| Legend |  | X Struck Note | — Held Note | O Optional: Struck or Held |  |

while the $(2m-2)$-gram and the $(2m-1)$-gram extend across at most $m$ quarter notes, the $(2m)$-gram and the $(2m+1)$-gram each extend across at most $(m+1)$ quarter notes. Examples of this phenomenon are shown in Table IV, while the counts of $m$-grams in Table II, having similar extension properties, are joined with dashes.

The synthesis for $m=1$ is random, obeying the (monogram) probability distribution. Any generated held note was interpreted as a continuation of the preceding struck note. A monogram hymn is given in Example 1 (right). The particular constraint used for this hymn is shown in Table V, along with constraints for the following examples. The low probability of accidentals (sharps and flats) in the sample, less than one per cent, resulted in the presence of only one accidental in Example 1: the *E* flat for the forty-third and forty-fourth notes. Despite this, however, the hymn is not easy to sing and contains unnatural intervals.

The digram hymn given as Example 2 exhibits several interesting irregularities. The constraint is a dotted constraint with two optional eighths (see Table V). Although the trigram *G F G* exists nowhere in the sample, this combination appears twice in the example as indicated by the brackets. In both cases, the optional cell contains the struck note *F* even though the held note $\overline{G}$ is much more likely to follow the *G*. Finally, nowhere in the sample does a *G* precede a *C* dotted half at the end of a phrase, nor does a second phrase begin on *F* after the first phrase ends on *C*. Indeed, none of these features are in keeping with the usual explicit rules of composition, but they are permitted by the inadequacy of the low order of the analysis-synthesis.

In the syntheses of orders 4 and 5, there is less roughness of the generated hymn. A tetragram hymn with no options in the metric constraint (Table V) is shown as Example 3. The problem of excessive range is one which was introduced implicitly by the naive method of transposition. Most of the original hymns had melodic lines with a range of about an octave, but the transpositions spread these ranges away from the normal vocal range.

The tetragram hymn illustrates a subtle manifestation of the first synthesis hindering effect, one that was not anticipated. In syntheses of intermediate order, there were long ascending or descending sequences each made up of a succession of the short ascents or descents so common in the sample. With higher-order procedures, these overlong sequences cannot occur.

Example 1 ($m=1$)

Example 2 ($m=2$)

Example 3 ($m=4$)

Example 4 ($m=6$)

Example 5 ($m=8$)

In Table III it is seen that the yield with $m=7$ represents a minimum for each constraint and that the yield with $m=6$ is quite near this minimum. A hexagram hymn generated with a basic quarter-note constraint is given in Example 4. This hymn demonstrates the existence of the "middle ground," and nowhere contains more than four consecutive quarter notes of any hymn in the sample. It shows the long-descent effect to some degree.

The yield for $m=8$ in Table III is appreciably greater than the yield for $m=7$. An examination of the octogram hymns reveals that a few of them are wholly identical with hymns in the sample. Several others have the entire first phrase of one hymn in the sample and the entire second phrase of another. An output hymn of order eight, which is an interesting composite of three hymns, is given as Example 5. The constraint is of the skeletal type and is shown in Table V. The seven-note section in brackets is common to two hymns at the segments and permits the changing of hymns in midphrase in this otherwise degenerate case. In a large number of output hymns in which such segmentation occurs, the transition preserves the absolute time coordinate of the original hymn. This was not true for orders lower than the eighth. Hence, at the eighth order, the third synthesis-hindering effect originally predicted has appeared. The sample is so redundant that the synthesis will not yield "original" utterances when carried out beyond the eighth order.

The skeletal constraint managed to produce a greater yield than the quarter note constraint for $m=8$. On the lower orders of synthesis ($m=4$ and 5), the former constraint permitted the hymn to run so astray that it could not meet a subsequent constraint. It hence produced fewer acceptable hymns than the quarter-note constraint. For $m=8$, however, the implicit structures revealed by the octogram analysis were strong enough to keep the synthesis within the framework of the sample, thus producing mostly hymns of which each half was exactly like half of a sample hymn. This is a further result of degeneracy.

## EXTENSION

The present experiments have permitted the identification and characterization of the limiting effects that apply to any generalized analysis and synthesis, any induction and subsequent deduction. Since these processes will become more and more important in the application of computing machinery to more delicate and sophisticated tasks, it is desirable to explore their characteristics and properties more fully. It is important to develop some measure of structural complexity in terms of information content, to develop better ways of characterizing the extent of limitation imposed by constraints, and to develop methods of describing the information content of a sample which is sufficient to permit synthesis of original members of some structure class.

Ideally, the present experiments would not have been needed. One would have described the information content of the sample and the extent of the constraints of the musical structure, and from this one could have predicted that first and second-order Markoff analysis-synthesis would have yielded sequences unacceptable as hymns and that eighth order analysis-synthesis would have yielded sample members.

The wide discrepancy between the ideal situation and that which currently prevails emphasizes the large amount of theoretical and experimental work which will have to be done before the inductive-deductive processes are well enough understood for general use in computing machine applications.

## ACKNOWLEDGMENT

The authors are indebted to their colleagues at the Harvard Computation Laboratory for advice and encouragement. Profs. K. E. Iverson, R. C. Minnick, and A. G. Oettinger gave especially valuable assistance. We are also indebted to Prof. H. H. Aiken for making available the Laboratory facilities for this investigation.