H. Kuensch, S. Geman, Ath. Kehagias.
"Hidden Markov Random Fields."

This paper has appeared in the journal:
Annals of Applied Probability, Vol.5, pp.577-602, 1995.

Hidden Markov Random Fields

Hans Künsch
Seminar fur Statistik
ETH Zentrum
CH-8092 Zurich
SWITZERLAND

Stuart Geman and Athanasios Kehagias
Division of Applied Mathematics
Brown University
Providence, RI 02912 USA

February, 1993

[1]Supported by Army Research Office contract DAAL03-92-G-0115 to the Center for Intelligent Control Systems, National Science Foundation grant DMS-8813699, and Office of Naval Research contract N00014-91-J-1021.

[2]Running head: Hidden Markov Random Fields

Abstract

A noninvertible function of a first order Markov process, or of a nearest- neighbor Markov random field, is called a hidden Markov model. Hidden Markov models are generally not Markovian. In fact, they may have complex and long range interactions, which is largely the reason for their utility. Applications include signal and image processing, speech recognition, and biological modeling. We show that hidden Markov models are dense among essentially all finite-state discrete-time stationary processes and finite-state lattice-based stationary random fields. This leads to a nearly universal parameterization of stationary processes and stationary random fields, and to a consistent non-parametric estimator. We show the results of attempts to fit simple speech and texture patterns.

1 Introduction

If $X = X_1, X_2, ...$ is a Markov process and $Y = Y_1, Y_2, ...$ is a deterministic or stochastic function of X, then Y is called a *hidden Markov model* (HMM), or sometimes a *hidden Markov process*. Usually, the dependency of Y_t on X is more-or-less local, as when $Y_t = f(X_t)$ for some function f, or $Y_t = g(X_t, X_{t+1}, \eta_t)$ for some function g and an iid process $\{\eta_t\}$, independent of X. In any case, Y itself is generally not Markov, and may in fact have a complicated dependency structure. Nevertheless, the *conditional* distribution of X given Y may remain simple, as in the above two examples where X given Y is still first-order Markov. The combination of a rich marginal structure

for Y and a simple posterior structure for X makes hidden Markov processes a common modeling tool.

Example: Filtering (cf. [33]). Although the general (nonlinear) filter problem falls within this framework, let us specialize to the linear case: X (known as the *state* process) is not only Markov, but satisfies a simple linear (stochastic) difference equation

$$X_{t+1} = aX_t + \omega_t$$

where $\{\omega_t\}$ is iid. The *observation* process Y is an HMM, linearly related to X, as in

$$Y_t = bX_t + \omega'_t$$

where $\{\omega'_t\}$ is another iid noise process, independent of $\{\omega_t\}$. The object is to estimate the state X_t from the observations $\{Y_s\}$, $s \in [0, T]$. This is termed *smoothing* if $0 \leq t < T$, *filtering* if $t = T$, and *prediction* if $t > T$. In any case, the fact that X given Y is still Markov is central to obtaining practical estimation formulas. Beyond this, linearity is exploited to derive efficient recursive estimators (e.g. the Kalman filter) for a host of "on-line" applications in tracking and control.

Example: Speech Recognition (See, for example, [1] and [41]). X is a Markov chain with finite (but very large) state space. In principle, the state of X_t represents all of the information relevant to predicting utterances of a speaker at times $\tau > t$. In practice, this information is modeled by representing, jointly, the word (and, sometimes, word pair), phoneme, and part of phoneme (e.g. beginning, middle, or end) being articulated at time t. The transition matrix for X is built hierarchically, by successively modeling the

variations in pronunciation of parts of phonemes, phonemes, and words, as well as (some of) the constraints and regularities in word sequences (syntax). Observations are of the acoustic signal, or some transformation or simplification, and are represented by Y. A stochastic model for Y_t given X_t is developed (or estimated more-or-less nonparametrically). The result is an HMM for the observable acoustic signal (or its transformation) Y, and the object is to estimate X (especially the word sequence) given Y. The posterior is Markov, which is fortunate since this simple dependency structure admits dynamic-programming-like computational tools for the calculatio (or at least approximation) of an optimal estimator for X, as well as for computing expectations of various sufficient statistics involved in the estimation of the model parameters. This HMM set up, or some of its variations, is the basis for the most successful speech recognition systems.

Example: Ion Channel Kinetics (See [3], [2], [23], and [35]). Nerve cells can propagate electrical activity without attenuation over long distances. Loss-less conduction involves an active process of opening and closing selective membrane ion channels, and thereby exchanging selected ions between inter and intra cellular spaces. Experiments can be devised to measure the changing conductance of one or a small number of channels in response to various chemical or electrical stimuli. These experiments reveal that ion channels typically move through only a few effective states, being, for example, simply "open" or "closed" with essentially no intermediate levels of conductance. The actual molecular basis for these measurable states is more complicated, and is often modeled as a Markov process with multiple states. The observable conductance is then a function of this process, through which,

for example, certain of the molecular states manifest themselves as an open channel, and others as a closed channel. Thus the observable conductance is an HMM. Purported mechanisms for channel kinetics can be tested by using observed channel conductances to infer the structure and transition probabilities of the (hidden) molecular Markov process. In these applications, the time parameter is generally continuous.

Example: Amino Acid Sequence Analysis. Hundreds or thousands of amino acids strung linearly together constitute a protein. Typically, there are only twenty distinct types of amino acids found, but there are of course a very large number of possible *sequences.* The particular sequence of amino acids that constitutes a protein is known as its "primary" structure. The determination of primary structure is known as sequencing, a process that has been increasingly automated; the result is a large existing data bank of primary structures. The *function* of a protein is largely determined by the folded three-dimensional (or "tertiary") structure that the amino acid chain assumes *in vivo.* Tertiary structure can sometimes be determined by experimental and imaging techniques, but the process is laborious and the number of sequenced proteins far exceeds the number of proteins with known tertiary structure. Hence, a fundamental problem in biology is the prediction of tertiary structure from primary structure.

One general approach is to search through sequences with known tertiary structures in order to find a "good match" to a sequence with unknown tertiary structure. Similar sequences tend to have similar structure, and in fact there are broad categories of structure that most proteins (or portions thereof) fall into. In an effort to exploit these structural categories, Krogh

et al. [37] build probabilistic models for amino acid sequences conditional on structural classes. These models are built up from known structure-sequence pairs, and then used to infer a likely structural class for a novel amino acid sequence. Thus, for example, a stochastic model is built for the sequence of amino acids constituting a typical globin (protein that transports oxygen and carbon dioxide). A new amino acid sequence can be evaluated under the globin model to determine its fit, and thereby to predict whether or not it will exhibit a globin-like tertiary structure. Preliminary tests have been highly successful.

The actual models constructed by Haussler et al. are HMM's, with the amino acids constituting the observables, and a Markov process, with carefully constructed state space and restricted transitions, constituting the hidden process. (A very similar approach is taken by Churchill [15] in constructing HMM's for the sequence of bases constituting a DNA molecule.) Transition probabilities are estimated from existing data bases, as are state-dependent distributions on the twenty available amino acids. Here again the conditional Markov structure of the unobserved (in fact, virtual) process is heavily exploited to develop computationally-feasible estimation and inference algorithms (involving various dynamic programming-like procedures).

Example: Texture Models. This is just a proposal, but it serves to introduce a generalization that will be a primary focus of our theoretical development. Consider a digitized image of a textured pattern such as cloth, wood, or sand. The image can be thought of as a realization of a stochastic process $\{Y_t\},\ t \in \Lambda = \{(i,j):\ 1 \le i \le N,\ 1 \le j \le M\}$, where $N = M = 512$ for example. Y_t is the grey level observed at picture element (or pixel) t. Many

authors (for example, [18], [32], and [20]) have proposed modeling $\{Y_t\}$, conditioned on the texture type and the imaging parameters (distance to camera, orientation, discretization, etc.), as a Markov random field. Since there is usually an organization to the texture that essentially rules out nearest-neighbor models, this approach demands that one either pick, more-or-less arbitrarily, a neighborhood structure, or attempt to estimate the neighborhood structure. In either case, there is then the requirement of choosing (or estimating) parameters that determine the associated clique functionals.

A different approach to obtaining the necessary structure would be to employ a *hidden* Markov random field, using a simple nearest-neighbor process for the underlying Markov structure. Thus $Y_t = f(X_t)$, f a fixed "hiding function," where X_t is a nearest-neighbor Markov random field. As in the one-dimensional examples discussed previously, Y will not generally be Markov, although the conditional distribution on X, given Y, is still a nearest-neighbor Markov random field. Is it possible to introduce sufficiently rich structure into the Y process to capture the regularity/variability of real textures through this mechanism? We will return to this shortly.

The last example, especially, raises the issue of generality: How general is the class of processes that can be well-approximated by a hidden Markov model? To be concrete, we shall restrict ourselves to nearest-neighbor processes (which is to say, first-order Markov when working in one dimension) and we will only allow instantaneous and deterministic "hiding" functions: $Y_t = f(X_t)$.[1] Furthermore, X_t (and hence also Y_t) will always have finite state

[1]In one dimension, many variations are popular: Y_t might depend, randomly or deterministically, on X_t or, simultaneously, on X_t and X_{t-1}. Restricting to *finite* state

space. So, for example, consider a stationary process $Z_t \in \{0,1\},\ t = 1, 2, ...,$ which we shall try to model (or "fit") with an HMM of the form $Y_t = f(X_t)$, where X_t is first-order Markov, $X_t \in \{0, 1, ...N\},\ f : \{0, 1, ...N\} \to \{0, 1\}$. By varying $N,\ f$, and the transition probability matrix for X, how close can we get (how similar to Z can we make Y)?

The answer depends very much on the measure of similarity. Ornstein and Weiss [38], for example, study related questions under a strong notion of similarity: Given two discrete-state stationary processes Y and Z, $d(Y, Z) \leq \epsilon$ if there exists a stationary process $\Psi = \{\Psi_t\} = \{(Y'_t, Z'_t)\}$ such that

1. Y' and Z' have the same distributions as Y and Z respectively, and
2. $P(Y'_1 \neq Z'_1) \leq \epsilon$.

The Ornstein-Weiss *distance*, d, between Y and Z is the *infimum* over all such ϵ. The results of Ornstein and Weiss indicate that the class of Z which can be arbitrarily well approximated by HMM's Y, relative to d, is highly restricted.

On the other hand, in terms of weak convergence, every stationary Z is the limit of a sequence of HMM's: There exists X^n, first order Markov on $\{0, 1, ...N_n\},\ f^n : \{0, 1, ...N_n\} \to \{0, 1\}$, such that $Y^n = f^n(X^n)$ converges weakly to Z as $n \to \infty$ (i.e., for every m, the distribution of $(Z_1, ...Z_m)$ is the limit of the distributions of $(Y^n_1, ...Y^n_m)$). This is fairly easy. Basically

spaces, it is not difficult to show that these four classes are *equivalent*, in the sense that the set of achievable distributions, for the observable process Y, is identical in each case (see [5], [34]). One constructs an explicit distribution-preserving transformation from an HMM of one type to an HMM of another type.

the idea is to define X^n taking values in $\{0,1\}^n$ ($X_t^n \in \{0,1\}^n$ for each $t = 1, 2, ...$) in such a way that X_{t+1}^n, given X_t^n, has the same distribution as $(Z_{t-n+2}, Z_{t-n+3}, ...Z_{t+1})$ given $(Z_{t-n+1}, Z_{t-n+2}, ...Z_t)$. Then Y_t^n is just the last component of X_t^n (see [38] and [34], and see [39] for versions of this for continuous-valued processes.)

The issue of approximating stationary processes by weak limits of HMM's is more complicated in higher dimensions. Let $S = \mathcal{Z}^d$ be the d-dimensional (discrete) square lattice. Let $Z = \{Z_t\}_{t \in S}$ be stationary with finite state space E ($Z_t \in E$, $\forall t \in S$). $X = \{X_t\}_{t \in S}$ is a nearest-neighbor Markov random field (MRF) if the distribution on X_t given $\{X_s\}_{s \neq t,\ s \in S}$ is the same as the distribution on X_t given $\{X_s\}_{s \in N_t}$, where N_t is the set of $2 \cdot d$ nearest neighbors of t (see [36]). When $d = 1$, this is equivalent to the usual first-order Markov property. Given stationary Z, can we choose an N, an X, and an f, such that X is a nearest-neighbor MRF, with values in $\{0, 1, ...N\}$, $f : \{0, 1, ...N\} \to E$, and the process Y defined by $Y_t = f(X_t)$ approximates Z? As we shall see shortly (§2), there always exists a sequence of these hidden nearest-neighbor Markov random fields that converges weakly to Z. (We use a similar idea as for one dimension: We choose as X_t^n the vector with components Z_s where s belongs to a block of pixels of size n around t. Actually, we will insist that our hidden process X^n be Gibbs—see §2, in addition to being Markov. This entails a modification to enforce strict positivity of the conditional probabilities, X_t^n given $\{X_s^n\}_{s \neq t}$.)

Given a stationary process $Z = \{Z_t\}_{t \in S}$, $S = \mathcal{Z}^d$, taking values in E ($!E! < \infty$), one way to actually build a model for Z would be to try and exploit the above-mentioned result about the (weak) density of HMM's:

Search for an N, a nearest-neighbor process $X = \{X_t\}_{t\in S}$, and a function $f : \{0, 1...N\} \to E$ such that $Y = \{Y_t\}_{t\in S}$, $Y_t = f(X_t)$, has distribution similar to Z. Actually, f can be fixed, *a priori*. For example, if $E = \{0, 1, ...M-1\}$, and $f(x) = x \bmod M$, then the collection of HMM's $Y_t = f(X_t)$, where X_t is a finite state ($\{0, 1, ...N\}$, some N) nearest-neighbor Markov random field on S, is weakly dense among all stationary Z (with state space $\{0, 1, ...M-1\}$) on S. Therefore, the construction of a model of this type amounts to choosing a suitable N and an associated process X. If $d = 1$, then X is determined by a transition probability matrix P, which requires specifying approximately N^2 parameters. If $d > 1$, then we can represent X as a Gibbs distribution (see §2), which will involve one pair-clique function for each dimension—roughly $d \cdot N^2$ parameters.

In §3 we address the problem of modeling Z by *estimating* these parameters via maximum likelihood (ML), or a closely-related methodology. We establish a kind of consistency result: Imagine that we are given a sequence of partial observations from a single realization of Z, of the form $\{Z_t\}_{t\in V_n}$, where $\{V_n\}_{n=1}^{\infty}$ is a sequence of increasing sublattices in $\mathcal{Z}^d$. The number N of states in the hidden process X amounts to a *regularization* or *smoothing* parameter, and, as is usual in nonparametric estimation, it will be necessary to relax the smoothing constraint as we accommodate more observations: $N = N_n \uparrow \infty$. We will present conditions under which a maximum likelihood (or closely-related) choice of the parameters of the X-process, under the hidden model $Y_t = f(X_t)$, guarantees a consistent estimation of Z, provided $N_n \uparrow \infty$ sufficiently slowly. Convergence is of a relative entropy (between Z and Y), and is almost sure with respect to the distribution on Z. Unfor-

tunately, we can offer no practical recipes for choosing N_n, or for actually calculating (global) ML estimates. Nevertheless, we performed some estimation experiments, fixing $N_n = N$ and $V_n = V$, involving acoustic signals from speech and simple binary textures; these are presented in §3 as well.

Related Work. We have already cited a few related papers. Additionally, several authors have addressed the problem of *identifiability:* given an HMM Y, describe the (generally large) class of Markov processes X that could, through a suitable hiding function f, generate the distribution of Y. Blackwell and Koopmans [10] seem to have been the first to address the problem. Their results were improved upon by Gilbert [28]. More recently, Ito et al. [30] obtained an essentially complete solution. Another related line of research has been the attempt to characterize, in terms of distributional properties, processes Y that are *exactly* functions of Markov chains. Dhamadhikari [19] gave some sufficient conditions, and Fredkin and Rice [22] gave some (rather severe and surprising) necessary conditions. A complete algebraic characterization is known, but it is not very manageable—see Chapter III of Rosenblatt [42]. Berbee and Bradley ([7], [11]) have constructed examples that show that even very rapidly mixing processes need not be HMM's. Brockett's calculations [12] indicate that good approximations of a stationary process by an HMM may require very large state spaces for the underlying Markov process, particularly when the stationary process has a nearly periodic covariance.

Concerning *estimation*, Baum and Petrie [5] established consistency of maximum likelihood estimation of an HMM when the state space of the (hidden) Markov process is known, and more recently Bickel and Ritov [8]

extended these results to include information about asymptotic distributions. The problem is, of course, harder when the hidden process (again, with known state space) is a Markov random *field*, but there has been progress here as well—see Comets and Gidas [16] and Frigessi and Piccioni [24]. The issue of how to actually compute maximum likelihood parameter estimators, both for hidden Markov processes as well as hidden Markov random fields, is discussed by Qian and Titterington [40], who suggest several variations on the E-M algorithm ([4], [6]), [17]), and Younes [43], who derives a stochastic gradient ascent algorithm. Finally, we mention the results of Ji [31], who studies nonparametric estimation of certain Gibbs fields. These are related to our estimation results, since our results amount to a recipe for nonparametrically estimating essentially arbitrary (stationary) random fields (see §3), although, unlike Ji, we give no information about rates of convergence.

2 Approximation

2.1 Notation and Preliminaries

As in §1, S will represent $\mathcal{Z}^d$, the d-dimensional discrete square lattice. Given any finite set E (such as the state space for either the hidden or observable process), the corresponding "configuration space" is $\Omega = E^S = \{x = \{x_t\}_{t\in S} : x_t \in E \ \forall t \in S\}$. The topology on Ω is, as usual, the product topology arising from the discrete topology on E. Similarly, if E' is another finite set, then $\Omega' = E'^S$, again with the product topology. Finally, if $V \subseteq S$, and $x \in \Omega$, then $x_V = \{x_t\}_{t\in V}$.

Gibbs measures (on Ω) are special cases of Markov random fields. We will show that the class of hidden finite-state first-order stationary Gibbs measures is weakly dense among finite-state stationary processes. In particular, this implies the result announced in §1, since first-order Gibbs measures are nearest-neighbor Markov random fields.

Gibbs measures arise from *potentials*. For our purposes we will use only shift-invariant and summable potentials. By this we mean a collection of functions $\Phi = \{\Phi_V\}_{V \subset S,\ V\ finite}$, such that

1. $\Phi_V :\ E^V \to R$,

2. $\Phi_{V+t} = \Phi_V\ \forall t \in S$,

3. $\sum_{V \ni 0} \sup_{x \in \Omega} |\Phi_V(x_V)| < \infty$.

A Gibbs measure with potential Φ is any probability measure μ on Ω such that, for any finite $V \subset S$ and $x \in \Omega$:

$$\mu[X_V = x_V | X_{V^C} = x_{V^C}] = \frac{1}{Z} \exp\{- \sum_{\substack{W \subset S \\ W \cap V \neq \emptyset}} \Phi_W(x_W)\}$$

where $V^C = S \setminus V$, and Z (which depends on V and x) normalizes the conditional distribution:

$$Z = \sum_{x_V} \exp\{- \sum_{\substack{W \subset S \\ W \cap V \neq \emptyset}} \Phi_W(x_W)\}.$$

Define $\delta V = \{t \in S \setminus V :\ \exists s \in V,\ W \subset S,\ |W| < \infty,\ \ni \Phi_W \neq 0\ and\ t, s \in W\}$, which is the boundary of V under the neighborhood relation induced

by Φ. Then $\mu[X_V = x_V | X_{V^C} = x_{V^C}]$ depends only on $x_{\delta V}$, so μ is an MRF relative to the neighborhood system

$$N_t = \{s \in S : \ s \neq t, \ s,t \in W \ some \ W \subset S, \ |W| < \infty, \ \Phi_W \neq 0\}.$$

In particular, if $\Phi_V = 0$ except for $V = \{t\}$ or $V = \{t, s\}$, where s and t are nearest neighbors in $\mathcal{Z}^d$ (so Φ is a "nearest-neighbor potential"), then μ is a nearest-neighbor Markov random field. A "stationary first-order Gibbs measure" is a stationary Gibbs measure with nearest-neighbor potential.

2.2 Statement of Result

Given E and E' finite, and a function $f : E' \to E$, denote by $\bar{f}$ the function from Ω' to Ω defined by

$$\bar{f}(\{x_t\}_{t \in S}) = \{f(x_t)\}_{t \in S}.$$

Given a measure ν on Ω', define $\mu = \nu \circ \bar{f}^{-1}$ on Ω by

$$\mu(A) = \nu(\bar{f}^{-1}(A)).$$

Now fix E and consider the following sets of probability measures:

$$\begin{aligned}
\mathcal{M}_s &= \{\mu \ on \ \Omega : \ \mu \ stationary\} \\
\mathcal{M}_g(E') &= \{\mu \ on \ \Omega' : \ \mu \ stationary \ first-order \ Gibbs\} \\
\mathcal{M}_h &= \{\mu \ on \ \Omega : \ \mu = \nu \circ \bar{f}^{-1}, \ for \ some \ E' \ finite, \ \nu \in \mathcal{M}_g(E'), \ f : E' \to E\}
\end{aligned}$$

$\mathcal{M}_h$ is the set of hidden finite-state first-order stationary Gibbs measures.

Theorem 2.2.1 *$\mathcal{M}_h$ is weakly dense in $\mathcal{M}_s$.*

Remark: In one dimension, $X = \{X_t\}_{t\in\{0,\pm1,\ldots\}}$ is first-order Markov with positive transition probabilities iff X is a first-order Gibbs. Hence, by the theorem, $\{Y : Y_t = f(X_t),\ some\ f\ and\ some\ X_t\ (finite-state)\ first-order\ Markov\ with\ positive\ transition\ probabilities\}$ is weakly dense among finite-state stationary processes. This special case is fairly easy to get (along the lines of the argument outlined in §1). Furthermore, in this case there are results about approximation in the sense of *relative entropy*—see [34].

2.3 Proof

The idea is essentially this: Gibbs measures are known to be dense in $\mathcal{M}_s$. Any Gibbs measure can be approximated by a Gibbs measure with potential having finite range: $\Phi = \{\Phi_V\}$, where $\Phi_V = 0$ whenever $diameter(V) > B$, for some bound B. Finally, *hidden* first-order Gibbs measures approximate Gibbs measures with finite-range potentials.

In general, there is more than one Gibbs measure with a given potential Φ ("phase transition"), and, even though Φ is shift-invariant, a Gibbs measure with potential Φ need not be stationary. Denote by $\mathcal{G}_s(\Phi)$ the set of all stationary Gibbs measures with potential Φ. Let $\mathcal{U}$ denote the set of potentials for which $\mathcal{G}_s(\Phi)$ is a singleton, and let

$$\mathcal{M}_u = \{\mu\ on\ \Omega :\ \{\mu\} = \mathcal{G}_s(\Phi)\ for\ some\ \Phi \in \mathcal{U}\}.$$

Then the following is known (Georgii [27], 16.40):

Proposition 2.3.1 *$\mathcal{M}_u$ is weakly dense in $\mathcal{M}_s$.*

Hence, it is sufficient to show that $\mathcal{M}_h$ is weakly dense in $\mathcal{M}_u$. The next step is to truncate Φ in order to have a finite range potential: Let

$$\Phi_V^N = \begin{cases} \Phi_V & \text{if } V \text{ is contained in } C_N + t \text{ for some } t \in S \\ \\ 0 & \text{otherwise} \end{cases}$$

where C_N is the cube $\{-N, -N+1, \ldots, N\}^d \subset S$. Then, in light of the following proposition, it will be sufficient to approximate for each N some member of $\mathcal{G}_s(\Phi^N)$ by a sequence in $\mathcal{M}_h$.

Proposition 2.3.2 *Suppose that $\Phi \in \mathcal{U}$ and that we have a sequence of potentials (Φ^N) such that $\sum_{V \ni 0} \sup_x |\Phi_V(x_V) - \Phi_V^N(x_V)| \longrightarrow 0$, as $N \to \infty$. Then for any sequence (μ_N) with $\mu_N \mu_{N_n} \to \mu$ weakly, where μ is the unique element of $\mathcal{G}_s(\Phi)$.*

Proof of Proposition 2.3.2. The hypothesis on (Φ^N) implies that the conditional probabilities

$$\Pi_V^N(x_V | x_{V^c}) := \frac{1}{Z_V^N} \exp(- \sum_{W \cap V \neq \emptyset} converge in the sup - norm to$$

$\Pi_V(x_V | x_{V^c}) := Z_V^{-1} \exp(-\sum_{W \cap V \neq \emptyset} \Phi_W(x_W)) for any V. Because \Omega$ is compact, we may assume that $\mu_N \to \nu$ weakly for some $\nu \in \mathcal{M}_s$. We have to show that $\nu \in \mathcal{G}_s(\Phi)$. First, we observe that it is enough to show

$$\int f d\nu = \int \Pi_V f d\nu \tag{1}$$

for any V and any f which depends only on x_V. (Here, $\Pi_V f(x)$ is defined as $\sum_{\xi \in E^V} \Pi_V(\xi | x_{V^c}) f(\xi)$.) Reason as follows: for any $V \subseteq W \subseteq S$, and f_1

depending only on x_V, and f_2 depending only on $x_{W\setminus V}$,

$$\Pi_W f_1 f_2 = \Pi_W f_2 \Pi_V f_1.$$

Now integrate both sides under ν measure, and apply (1):

$$\int f_2 f_1 d\nu = \int f_2 \Pi_V f_1 d\nu$$

which implies that ν has the right conditional probabilities (i.e. Π_V) on E^V, which, in turn, implies that $\nu \in \mathcal{G}_s(\Phi)$. So, we need to prove (1).

Because $\Pi_V^N \to \Pi_V$, we have $\sup_x !\Pi_V^N f(x) - \Pi_V f(x)! \to 0$, so

$$!\int f d\nu - \int \Pi_V f d\nu! \leq !\int f d\nu - \int f d\mu_N !+ !\int f d\mu_N if$$

$Nislargeenough.Thefirsttermgoestozerobecause\mu_N \to \nu$, and the second term is zero because $\mu_N \in \mathcal{G}(\Phi^N)$. q.e.d.

Hence the theorem follows from

Proposition 2.3.3 *Suppose Φ is a potential whose range is contained in some C_N. Then there is a sequence $\{\nu_\beta\}$ of stationary first-order Gibbs measures with state space $E' = E^{C_N}$ and a function $f : E' \to E$ such that*

Proof of Proposition 2.3.3. Since N is fixed we write C instead of C_N. If $\{x_t\} \in \Omega = E^S$, then we define $\{y_t\} \in \Omega' = E'^S$ by

$$y_t = x_{C+t}. \tag{2}$$

We index the components of y_t by $r \in C$, rather than choosing an arbitrary enumeration. So the r-th component of y_t, $y_{t,r}$, is equal to x_{t+r}. We will use the symbol y always to denote an element of Ω' of the form given by

(2). An arbitrary element of Ω' will be written as z or ξ. An equivalent way to express that y is of the form (2), is via the following compatibility constraints:

$$y_{t,r} = y_{s,r+t-s} \ \forall r, t, s \text{ with } r \in C \text{ and } r + t - s \in C. \tag{3}$$

Moreover the following apparently weaker form is also equivalent to (**??**) and (2):

$$\begin{aligned} y_{t,r} = y_{s,r+t-s} \ \forall r, t, s \text{ with} & \, t - s = 1, \ r \in C \\ \text{and} \quad & r + t - s \in C. \end{aligned} \tag{4}$$

This can be seen by connecting t and s through a chain $t = t_0, t_1, \ldots, \ t_n = s$ such that $t_{i+1} - t_i = 1$ and $r + t_{i+1} - t_i \in C$.

We now turn to the definition of the approximating potential Φ^β (which will define a first-order Gibbs measure on Ω'). To motivate this definition, note that $\{y_t\}$ is a sample from a first-order Markov random field if $\{x_t\}$ is a sample from some $\mu \in \mathcal{G}_s(\Phi)$. Howe because of the hard constraints (4). We change these hard constraints into soft ones by introducing potentials with value β for each of the constraints (4) which is violated. By letting β tend to infinity we hope to recover (4) and at the same time to obtain the right distribution on the configurations satisfying (4). We show that this is indeed the case.

We seek to approximate the potential Φ. Without limitation of generality we may assume:

$$\Phi_V \equiv 0 \text{ if } V \neq t + C, \text{ some } \ t \in S.$$

Then we define Φ^β as follows

$$\Phi^\beta_{\{t\}}(z_t) \qquad = \Phi_C(z_t)$$

$$\Phi^\beta_{\{t,s\}}(z_t, z_s) \quad = \beta\Psi(z_t, z_s) := \beta \sum_{r\in C^{r,}_{\text{and}}} 1_{[z_{t,r}\neq z_{s,r+t-s}]}$$

$$\text{if } t-s=1$$

$$\Phi^\beta_V \quad = \quad 0 \quad \text{otherwise}$$

We denote by π^β_V the conditional probabilities associated with Φ^β and choose some $\nu_\beta \in \mathcal{G}_s(\Phi^a ssumethat\nu_\beta \to \nu$ weakly as $\beta \to \infty$. We $z_{t,0}$. The following lemma is the key.

Lemma 2.3.1 *For arbitrary V consider the event $A = \{z \in \Omega' ! \Psi(z_t, z_s) = 0 \ \ \forall t, s \in V, t-s=1\}$. Then $\nu(A) = 1$, i.e. the compatibility constraints are fulfilled a.s.*

We defer the proof of this lemma to the end and consider the conditional distribution of $z_{0,0}$ given $z_{s,0}$ $s \neq 0$, under ν. W is the size of the box which contains the potential Φ) and put

$$\Lambda = C_K \backslash C_N$$

$$\bar{\Lambda} = C_{K+N} \backslash \{0\}.$$

We want to show that

$$\nu(z_{0,0} = x_0 ! z_{s,0} = x_s, s \in \bar{\Lambda}) \propto \exp\{-\sum_{t\in C} \Phi \tag{5}$$

Note that knowing $x_s, s \in \overline{\Lambda}$ is the same as knowing $y_s, s \in \Lambda$. Hence by Lemma 2.3.1 above

$$\nu(z_{s,0} = x_s, s \in \overline{\Lambda}) = \nu(z_s = y_s, s \in \Lambda).$$

Moreover if this probability is positive, then by weak convergence of ν_β to ν

$$\nu(z_{0,0} = x_0 ! z_s = y_s, s \in \Lambda) = \lim_{\beta \to \infty} \nu_\beta(z_{0,0} = x_0 ! z_s = y_s, s \in \Lambda)$$

$$= \sum_{\substack{z_C \\ z_{0,0} = x_0}} \lim_\beta \Pi_C^\beta(z_C ! y_\Lambda)$$

But by the definition of Φ^β we have with $x_0 = z_{0,0}$

$$\Pi_C^\beta(z_C ! y_\Lambda) = Z^\beta(y_\Lambda)^{-1} \exp\Big\{ - \sum_{t \in C} \Phi^\beta_{\{t\}}(z_t)$$

$$- \sum_{\substack{s,t \in C \\ s - t = 1}} \Phi^\beta_{\{s,t\}}(z_s, z_t) - \sum_{\substack{s \in C \\ t \in \Lambda \\ s - A s}}$$

$\beta \to \infty$ this converges to

It remains to prove Lemma 2.3.1.

Proof of Lemma 2.3.1. For any Λ, define the energy in Λ given the boundary conditions by

$$H^\beta_\Lambda(z) = \sum_{W \cap \Lambda \neq \emptyset} \Phi^\beta_W(z_W).$$

Choose Λ such that $V + 2C = V + C + C \subset \Lambda$. We will show that we can modify an arbitrary configuration $z \notin A$ to one which belongs to A, has

the same boundary conditions with respect to Λ, and whose energy in Λ is smaller by an amount β. Hence

$$\Pi_\Lambda^\beta[A!z_{S\setminus\Lambda}] \to 1 \text{ as } \beta uniformlyin$$

$z_{S\setminus\Lambda}$. By integrating over the boundary conditions we thus obtain

$$\nu_\beta[A] \to 1 \text{ as } \beta \to \infty.$$

The modification mentioned above goes as follows: Let $g : \Omega' \to \Omega'$ be defined by

$$g(z)_{s,r} = z_{s+r,0} \text{ if } r \in C \text{ and } s+$$
$$g(z)_{s,r} = z_{s,r} \text{ otherwise}$$

then it is clear that $g(z)_s = z_s$ if $s \notin V + 2C$, and thus the boundary condition does not change. Also, it is seen easily that $g(z)$ belongs to A. So let us compare the energies $H_\Lambda^\beta(z)$ and $H_\Lambda^\beta(g(z))$ for $z \notin A$:

$$\sum_{t\in\Lambda} \Phi_{\{t\}}^\beta(g(z)_t) \le \sum_{t\in\Lambda} \Phi$$

where $\delta = \sup \Phi_C(x_C) - \inf \Phi_C(x_C) because g(z)_t = z_t$ if $t \notin V + 2C$. Moreover, because $z \notin A$ and $g(z) \in A$

$$\sum_{\substack{t,s\in V\\ t-s=1}} \Psi(z_t, z_s) \ge 1, \sum_{\substack{t,s\in V\\ t-s=1}} \Psi(g(z)_t, g(z)_s) = 0$$

Now take a $t \notin V$ and s arbitrary with $t - s = 1$. If $t + r \in V + C$ and $r + t - s \in C$, then $g(z)_{t,r} = g(z)_{s,r+t-s}$, and if $tthen g(z)_{t,r} = z_{t,r}$ and $g(z)_{s,r+t-s} = z_{s,r+t-s}$. Hence $\Psi(g(z) Together we obtain H_\Lambda^\beta(g(z)) \le H_\Lambda^\beta(z) - \beta + !V + 2C!\delta. This completes the proof of the lemma.$ *q.e.d.*

3 Estimation

The approximation result of §2 suggests modeling stationary processes with hidden nearest-neighbor MRF's, or simply (hidden) first-order Markov processes in the one-dimensional case. A single sample path from an ergodic stationary process should be sufficient to determine the parameters for an approximation of this type, and this is confirmed, roughly speaking, by our consistency results: Given a sequence of observations from a single sample of an ergodic stationary process, we use Grenander's method of sieves [29] to construct a sequence of hidden first-order processes with distributions converging (in the sense of relative entropy) to the stationary process.

Actually, we will need to restrict the class of stationary processes, somewhat when working in one dimension, and somewhat more when working in higher dimensions; see §3.1.

Two sets of experiments were performed. Data from speech signals and textures were used to fit hidden nearest-neighbor processes, via maximum likelihood, and the resulting models were sampled and compared to the original data. See §3.2.

3.1 Consistency

There are two theorems, for random processes ($S = \mathcal{Z}$) and random fields ($S = \mathcal{Z}^d,\ d > 1$), respectively. The two proofs follow the same general plan, which we will present, in brief outline, for the one-dimensional ($S = \mathcal{Z}$) case. The full details are available in a technical report; see [26].

We imagine observing a stationary process Z with state space $E =$

$\{0, 1, ...M-1\}$, for some (finite) $M > 1$. Let μ_o be the (unknown) distribution, or law, of Z. Following the notation of §2, the process Z is to be approximated by a hidden process $Y = \bar{f}(X)$, where X is nearest-neighbor with state space $E' = \{0, 1, ...N\}$. Henceforth, the hiding function f (and consequently $\bar{f}$ as well) is fixed: $f(x) = x \bmod M$. Specializing to the one-dimensional problem ($S = \mathcal{Z}$), X is first-order Markov, and we will adopt the standard representation in terms of transition probability matrices rather than using potentials and the Gibbs representation. Let

$$\begin{aligned} \mathcal{M}_N &= \{m = \{m_{ij}\}_{i,j=0}^{N} : \ m \ trans.\ prob.\ matrix, \\ &\qquad and\ m_{ij} \geq e^{-N} \ \forall\ 0 \leq i, j \leq N\}. \end{aligned}$$

N will serve as a "regularization" or "smoothing" parameter, and will eventually be tied to the number n of observations, $Z_0 = z_0,\ Z_1 = z_1, ...Z_n = z_n$, through an increasing function. For any $m \in \mathcal{M}_N$, denote by μ_m the distribution of the hidden Markov process $Y = \{Y_t\},\ Y_t = f(X_t)$, where $\{X_t\}$ is the unique stationar Markov process with transition matrix m. The results of §2 suggest that μ_o can be approximated by a distribution μ_m, for suitable m and large enough N. Having observed $Z_0 = z_0,\ Z_1 = z_1, ...Z_n = z_n$, we denote by $ML_{N,n}$ the set of maximum likelihood matrices from within $\mathcal{M}_N$:

$$ML_{N,n} = ML_{N,n}(z) = \{m \in \mathcal{M}_N : \mu_m(z_0, z_1, \ldots z_n) = \sup_{q \in \mathcal{M}_N} \mu_q(z_0, z_1, \ldots z_n)\}.$$

(In general, $ML_{N,n}$ has more than one element. In any case, it is never empty: $\mathcal{M}_N$ is compact and μ_q is continuous in q.) Under an additional condition

on Z, there exists a sequence N_n such that the set of HMM's associated with $ML_{N_n,n}$ is consistent for μ_o:

Theorem 3.1.1 *Let $\{Z_t\}_{t=-\infty}^{\infty}$ be a stationary ergodic process with finite state space, $Z_t \in \{0, 1, \ldots M-1\}$, $M < \infty$, and distribution function μ_o. If $\exists\ \delta > 0\ \ni\ \mu_o(z_0|z_1, \ldots z_{-t}) \geq \delta\ \forall\ t,\ (z_0, \ldots z_{-t}) \in \{0, 1, \ldots M-1\}^{t+1}$, then for all $N_n \uparrow \infty$ sufficiently slowly*

$$\sup_{m \in ML_{N_n,n}} \int \log \frac{\mu_o(z_0!z_{-1}, z_{-2}, \ldots)}{\mu_m(z_0!z_{-1}, z_{-2}, \ldots)} d\mu_o(z) \to 0\ a.s.\ (\mu_o)$$

Remarks.

1. More precisely, there exists a sequence $N_n \uparrow \infty$ such that the assertion holds for all sequences $N'_n \uparrow \infty$ satisfying $N'_n \leq N_n\ \forall n$.

2. Unfortunately, $N_n = N_n(\mu_o)$; roughly speaking, $\{Z_t\}$ can yield information arbitrarily slowly.

3. There is nothing special about the regularization $m_{ij} \geq e^{-N}$. If instead, $m_{ij} \geq g(N)$, where $g(N) \downarrow 0$, then there will be a relationship between $g(N)$ and N_n such that the *faster* $g(N) \downarrow 0$ the *slower* $N_n \uparrow \infty$, in order to insure consistency.

The corresponding result for Z on $S = \mathcal{Z}^d,\ d > 1$, is somewhat more complicated, even to state. First, we make the additional assumption that μ_o, the distribution of Z, is in fact a (ergodic and stationary) Gibbs measure (see §2). In other words, we shall assume that μ_o is a measure on $\{0, 1, \ldots M-1\}^S$ satisfying

i. μ_o is stationary and ergodic,

ii. for every finite $V \subset S$

$$\mu_o[Z_V = z_V | Z_{V^c} = z_{V^c}] \propto \exp\{- \sum_{\substack{W \subset S \\ W \cap V \neq \emptyset}} \Phi_W(z_W)\}$$

where $\Phi = \{\Phi_V\}$ $V \subset S$, finite, is a (shift-invariant, summable) *potential*, as defined in §2.[2]

We can no longer index the approximating measures by transition probabilities; instead we replace $\mathcal{M}_N$ by a set of "regularized" potentials $\mathcal{P}_N$:

$$\begin{aligned} \mathcal{P}_N \;=\; & \{\textit{nearest neighbor potentials } \Psi = (\Psi_0, \Psi_1, \ldots \Psi_d) \\ & \textit{on } \{0, 1, \ldots N\} \textit{ with bounds } !\Psi_0(k)! \leq N, \\ & !\Psi_i(j,k)! \leq N,\ 1 \leq i \leq d\}, \end{aligned}$$

with the understanding that $\Psi_0 : \{0, 1, \ldots N\} \to R$ is the one-point potential $\Phi_{\{t\}}$ and $\Psi_i : \{0, 1, \ldots N\}^2 \to R$ are the pair potentials $\Phi_{\{t,t+e_i\}}$ $(1 \leq i \leq d)$, where e_i is the vector with $d-1$ zeroes and a single one at the i'th component. All other potentials Φ_V are identically zero. A Gibbs measure ν with potential $\Psi \in \mathcal{P}_N$ is then defined as in §2.1. Evidently, ν is then nearest-neighbor Markov.

[2]Obviously ii implies that the conditional distributions are bounded from below as was required already in Theorem 3.1.1. In addition ii also implies that the conditional distributions are continuous in z. As a converse, boundedness and continuity of the conditonal distributions imply ii (Georgii [27], 2.30). Furthermore, by Proposition 2.3.1 above we also know that the set of μ_o's satisying i and ii is weakly dense in the set of stationary measures since uniqueness of a Gibbs measure implies that it is stationary and ergodic (Georgii [27], 5.11 and 14.15).

For any $\Psi \in \mathcal{P}_N$, let $\mathcal{G}_s(\Psi)$ be the set of stationary Gibbs measures with potential Ψ. $\mathcal{G}_s(\Psi)$ is always nonempty, but may contain more than one measure. Therefore, there is a *set* of hidden Gibbs measures associated, through f, with each $\Psi \in \mathcal{P}_N$; we denote this set by $\mathcal{H}(\Psi)$:

$$\mathcal{H}(\Psi) = \{\mathcal{L}(Y) : Y_t = f(X_t),\ \mathcal{L}(X) \in \mathcal{G}_s(\Psi)\}$$

where $\mathcal{L}(\cdot)$ is the distribution (or law) of a process.

Now suppose that we observe $Z_V = z_V$, where $\mathcal{L}(Z) = \mu_o$ and $V \subset S$ finite. The idea is to choose a maximum likelihood potential Ψ from within $\mathcal{P}_N$, in other words to choose $\Psi \in \mathcal{P}_N$ in such a way that the associated hidden Gibbs measures assign maximum probability (likelihood) to z_V. Unfortunately, given a candidate potential $\Psi \in \mathcal{P}_N$, the likelihood of z_V under the hidden Gibbs model associated with Ψ is not necessarily well-defined; different elements of $\mathcal{H}(\Psi)$ may assign different likelihoods to z_V. Furthermore, even when $\mathcal{G}_s(\Psi)$ contains only one measure, the actual calculation of the probability of z_V under the associated hidden measure is intractable. For these reasons we will employ the following modification of the likelihood:

Fix, once and for all, a configuration $x \in \Omega = \{0, 1, \ldots M-1\}^S$. For any $\Psi \in \mathcal{P}_N$, define the (conditional) log-likelihood

$$L_V(\Psi, z_V) = \log\{\sum_{\substack{\xi_V:\ \xi_t \in f^{-1}(z_t) \\ t \in V}} \mu[\xi_V ! x_{V^C}]\},$$

where $\mu \in \mathcal{G}_s(\Psi)$. L is well defined: it is independent of which $\mu \in \mathcal{G}_s(\Psi)$ we choose. Furthermore, L depends only on $x_{\delta V}$, where δV is the boundary of V under the nearest-neighbor system in $\mathcal{Z}^d$.

Finally, define $M_{N,V}$ to be the set of maximum likelihood potentials within $\mathcal{P}_N$:

$$M_{N,V} = M_{N,V}(z) = \{\Psi \in \mathcal{P}_N : \ L_V(\Psi, z_V) = \sup_{\Phi \in \mathcal{P}_N} L_V(\Phi, z_V)\},$$

and for any two stationary probability measures μ and ν, define $h(\mu, \nu)$ to be the *specific relative entropy* (see e.g. Georgii [27]):

$$h(\mu, \nu) = \liminf_{V \uparrow S} \frac{1}{!V!} E_\mu[\log \frac{\mu[x_V]}{\nu[x_V]}].$$

By Jensen's inequality, $h(\mu, \nu) \geq 0$, and $h(\mu, \nu) = 0$ if $\mu = \nu$. Conversely, if ν is Gibbs with summable potential and $h(\mu, \nu) = 0$, then also μ is Gibbs with the same potential as ν (Georgii [27], 15.37).

Theorem 3.1.2 *Let μ_o be an ergodic stationary Gibbs measure on $\{0, 1, \ldots M-1\}^S$, $S = \mathcal{Z}^d$, and let $\{V_n\}$ be an increasing sequence of finite subsets of S, such that $S = \bigcup_{n=1}^{\infty} V_n$ and $!\delta V!/!V! \to 0$. For all sequences $N_n \uparrow \infty$ sufficiently slowly*

$$\sup_{\Psi \in ML_{N_n, V_n}} \sup_{\mu \in \mathcal{H}(\Psi)} h(\mu_o, \mu) \to 0 \quad a.s.\ (\mu_o)$$

Remarks.

1. For $d = 1$ the claim of Theorem 3.1.2 is the same as of Theorem 3.1.1, i.e.

   $$h(\mu_o, \mu) = \int \log(\mu_o(z_0!z_{-1}, \ldots)/\mu(z_0!z_{-1}, \ldots)) d\mu_o(z)$$

   for the law μ of any hidden Markov process. This is easily seen by an argument following the Shannon-McMillan-Breiman theorem since for any such μ

   $$\mu(z_0!z_{-1}, \ldots, z_{-k}) \longrightarrow \mu(z_0!z_{-1}, \ldots)$$

uniformly in z (cf. [26]).

2. Unfortunately we do not know whether we have convergence also with respect to the weak topology of measures.

Proof (outline). The approach is substantially the same for both consistency results. Let us consider the case of μ_o defined on $\{0, 1, \ldots M-1\}^{\mathcal{Z}}$ (i.e. dimension one), and go through a brief outline of the proof. (The details, for both the consistency theorems, are available through the technical report [26].)

The proof is based upon two lemmas. The first is a kind of uniform law of large numbers for the probabilities $\mu_m,\ m \in \mathcal{M}_n$, reminiscent of the Shannon-McMillan-Breiman Theorem (cf. Billingsley [9]):

Lemma 3.1.1

$$\lim_{n\to\infty} \sup_{m\in\mathcal{M}_{N_n}} !\frac{1}{n}\log \mu_m(z_0, z_1, \ldots z_n) - \int \log \mu_m(z_0!z_{-1}, z_{-2}, \ldots) d\mu_o(z)! = 0 \ \ a.s.\ (\mu_o)$$

for all $N_n \uparrow \infty$ sufficiently slowly.

The second lemma insures that there is *some* sequence $m_N \in \mathcal{M}_N$ such that μ_{m_N} approaches μ_o:

Lemma 3.1.2 *There exists a sequence of matrices $m_N \in \mathcal{M}_N$ such that*

$$\lim_{N\to\infty} \int \log \mu_{m_N}(z_0!z_{-1}, z_{-2}, \ldots) d\mu_o(z) = \int \log \mu_o(z_0!z_{-1}, z_{-2}, \ldots) d\mu_o(z).$$

(The proof of lemma 3.1.2 is by construction.)

Now assume that the lemmas are true. By Jensen's inequality,

$$\int \log\ \mu_m(z_0!z_{-1}, \ldots) d\mu_o(z) \leq \int \log\ \mu_o(z_0!z_{-1}$$

for all N and $m \in M_N$, so it is enough to show that

$$\liminf_{n\to\infty} \inf_{m\in ML_{N_n,n}} \int \log\ \mu_m(z_0!z_{-1},\ldots)d\mu_o(z) \geq \int\ \mu_o(z_0!z_{-1},\ldots)d\mu_o(z)\ a.s.$$

By application of the lemmas:

$$\liminf_{n\to\infty} \inf_{m\in ML_{N_n,n}} \int \log\ \mu_m(z_0!z_{-1},\ldots)d\mu_o(z)$$

$$\begin{aligned}
&= \liminf_{n\to\infty} \inf_{m\in ML_{N_n,n}} \{(\int \log\ \mu_m(z_0!z_{-1},\ldots)d\mu_o(z) - \frac{1}{n}\log\ \mu_m(z_0,z_1,\ldots z_n)) \\
&\quad + \frac{1}{n}\log\ \mu_m(z_0,z_1,\ldots z_n)\} \\
&\geq \liminf_{n\to\infty} \inf_{m\in ML_{N_n,n}} \{\frac{1}{n}\log\ \mu_m(z_0,z_1,\ldots z_n) - ! \frac{1}{n}\log\ \mu_m(z_0,z_1,\ldots z_n) \\
&\quad - \int \log\ \mu_m(z_0!z_{-1},\ldots)d\mu_o(z)!\} \\
&= \liminf_{n\to\infty} \inf_{m\in ML_{N_n,n}} \frac{1}{n}\log\ \mu_m(z_0,z_1,\ldots z_n)\ \ (a.s.,\ by\ lemma\ 3.1.1) \\
&\geq \liminf_{n\to\infty} \frac{1}{n}\log\ \mu_{m_{N_n}}(z_0,z_1,\ldots z_n) \\
&= \liminf_{n\to\infty} \int \log\ \mu_{m_{N_n}}(z_0!z_{-1},\ldots)d\mu_o(z)\ \ (again,\ a.s.,\ by\ lemma\ 3.1.1) \\
&= \int \log\ \mu_o(z_0!z_{-1},\ldots)d\mu_o(z)\ \ (by\ lemma\ 3.1.2)
\end{aligned}$$

3.2 Experiments

Consistency is reassuring, but it tells us too little about performance on real (finite) data. We have therefore run some simple experiments in order to asses the "finite sample" promise of the proposed models. There were two kinds of experiments: one-dimensional estimation experiments from speech waveforms and two-dimensional estimation experiments from simple binary textures.

3.2.1 Speech Waveforms

Segments of two phonemes were extracted from a single utterance of the word "one"—see Figure 1. The acoustic signal was sampled at xx kilohertz (one sample every xx milliseconds) and 256 amplitude levels, although each amplitude was later rounded to one of eight equally-spaced values. Panel A shows a xx millisecond segment from the phoneme [ah] which follows the initial [ou] and precedes the final [nn] in the pronunciation of "one." There are 200 data points, each having one of eight values. Panel B shows an analogous segment from the final [nn] of the same utterance. The nearly-periodic waveforms are characteristic of so-called voiced phonemes, and derive ultimately from more-or-less periodic oscillations of the vocal chords.

We treated each signal as a sample, $z_1, z_2, ...z_{200}$, from an eight-valued stationary process, which we attempted to fit with a series of hidden Markov models of increasing size. In each case, we employed the "hiding function" $f(x) = 1 + x \bmod 8$, and computed approximate maximum-likelihood $N \times N$ transition probability matrices, for $N = 10,\ 20,\ 30,\ 40,\ 50$, and 60. Maximum likelihood computations were made via the Baum re-estimation formula ([4], [6]), which is an instance of the EM procedure ([17]). Estimates were only approximately maximum likelihood since this is an iterative hill-climbing algorithm; it can approach a local maximum and, as a practical matter, it must be terminated short of convergence. We began each run (one run for each value of N) with a randomly-generated transition probability matrix, and continued until there were only negligible changes in the transition matrix.

The results are most easily judged by viewing samples from the resulting HMM's. Figures 2A (for the [ah] sequence) and 2B (for the [nn] sequence) show random samples from $Y_t = f(X_t^N)$, where $\{X_t^N\}_{t=1}^{200}$ is first-order Markov on $\{0, 1, ...N-1\}$ with the estimated $N \times N$ transition probability matrix, and $X_1^N = 1$ ($N = 10,\ 20,\ 30,\ 40,\ 50$, and 60). In both sets of experiments, the results appear to deteriorate at the largest values of N. Since there are only 200 (highly correlated) samples, the likely explanation is the familiar problem of over-fitting, although there may be computational problems (perhaps related to local maxima) with the iteration procedure. In any case, it would be interesting to perform similar experiments with larger data sets; essentially infinite amounts of data are easily available.

It may also be interesting to splice together such signals, as a novel approach to speech synthesis. In this regard, one would need to fit, as well, the non-stationary speech units associated with various consonants. Because we are after a signal of only finite duration, it is not impossible, and perhaps not unreasonable to speculate, that exactly the same models would be effective for fitting consonants.

An obvious alternative approach would be to fit each signal with an N'th-order Markov process. However, even at the modest eight-level discretization used in our experiments, this would involve estimating $7 \cdot 8^N$ parameters, which evidently places a severe restriction on the process order. It may be true, in contrast, that the hidden process provides an efficient coding of the nearly-periodic structure by dedicating single or multiple states to positions within the cycle, although we have performed no systematic experiments to test this conjecture.

3.2.2 Binary Textures

The experiments with two-dimensional processes were more difficult and less successful. We adopted the modest goal of fitting some simple binary textures. These were derived from real textures, borrowed from the well-used Brodatz collection ([13]), by simply thresholding grey-level pictures. A suitable threshold produces substantial islands of "ones" positioned among a sea of "zeros." The shape and pattern of the islands, of course, depends upon the texture. Figure 3 has two examples, straw and paper. In each, there are $60 \times 80 = 4800$ pixels; "ones" are depicted with dots and "zeros" with stars.

Following our approach to the speech data, we viewed these images as samples from stationary (spatial) processes, and attempted to fit these processes with N-state hidden Markov models. Specifically, we employed the hiding function $f(x) = x \bmod 2$, and a four-nearest-neighbor Gibbs representation for the hidden process, $X_t,\ t \in S = \{(i,j) : 1 \le i \le 60, 1 \le j \le 80\}$. In both experiments, N was fixed at 10, so that $X_t \in \{0, 1, ...9\}$.

For each texture we fit two matrices $\alpha^h = \{\alpha^h_{kl}\}$ and $\alpha^v = \{\alpha^v_{kl}\}$, where $0 \le k, l \le 9$, and h stands for "horizontal" and v for "vertical." These matrices represent the Gibbs potential for X, as follows:

$$\Pi(X_{i,j} = k | X_{i-1,j} = l_1, X_{i+1,j} = l_2, X_{i,j-1} = l_3, X_{i,j+1} =$$

$$\propto \exp -\{\alpha^v_{l_1 k} + \alpha^v_{k l_2} + \alpha^h_{l_3 k} + \alpha^h_{k l_4}\},$$

except that terms are dropped when they reference outside of the 60×80 array ("free boundary conditions"). Given a sample $z = \{z_t\}_{t \in S}$, the partial derivative with respect to α^h_{kl} $(0 \le k, l \le 9)$ of the log-likelihood of z, under

the model $\{f(X_t)\}_{t\in S}$, is

$$E[N^h_{kl}] - E[N^h_{kl}|f(X_t) = z_t, t \in S] \tag{6}$$

where

$$N^h_{kl} = \#\{(i,j) : 1 \le i \le 60, 1 \le j \le 79, X_{i,j} = k, X_{i,j+1} = l\}$$

(a "sufficient statistic"). An analogous expression governs partial derivatives with respect to the components of α^v. One way to estimate the matrices α^h and α^v is via a discrete gradient ascent: compute (6) at the "current" parameter values, take a small step in the direction of the gradient, recompute (6), and so on. Unfortunately, the computation of (6) is notoriously difficult. We resorted to Monte Carlo methods (cf. Metropolis et al. [**?**]xxx) and Besag and Green [**?**]xx)), repeatedly using the Gibbs Sampler to estimate both expectations.

The approach is unsatisfactory. It is slow and it is difficult to judge convergence, both within an iteration (computation of the expectations) and overall (when to stop?). There have been many suggestions for improving the efficiency of the calculations—see for example Younes ([43]) and Qian and Titterington ([40]). We experimented with a variety of alternatives, without much success. In the end we settled on the approach outlined above, which we view as decidedly brute-force and last-resort.

Having estimated potential functions (α^h and α^v) for both the (binarized) straw and paper textures, we drew samples from the corresponding Gibbs distributions—again, via the Gibbs Sampler. The results, viewed through the hiding function f, are shown in Figure 4.

As with the problem of synthesis in speech, texture synthesis is made intriguing by the availability of unlimited amounts of data. Despite this favorable circumstance, there are as of yet no fully satisfactory solutions, especially if one wants to render samples at arbitrary angles and resolution. We have offered a solution, *in principle*: nearest-neighbor HMM's are dense and can be estimated. Evidently, however, the approach is a long way from being practical. In any case, others have already made good progress— we would cite [14], [25], [18], [21], [32], and [20], for some state-of-the-art work on texture estimation and synthesis.

References

[1] L.R. Bahl, F. Jelinek, and R.L. Mercer. A maximum likelihood approach to continuous speech recognition. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 5:179–190, 1983.

[2] F. Ball and M. Sansom. Aggregated Markov processes, incorporating time interval omission. *Adv. Appl. Prob.*, 20:546–572, 1988.

[3] F.G. Ball and J.A.Rice. Stochastic models for ion channels: introduction and bibliography. *Math. Biosciences*, 112:189–206, 1992.

[4] L.E. Baum and J.A. Eagon. An inequality with applications to statistical estimation for probabilistic functions of Markov processes and to a model for ecology. *Bull. AMS*, 73:360–363, 1967.

[5] L.E. Baum and T. Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *Ann. Math. Statist.*, 37:1554–1563, 1966.

[6] L.E. Baum, T. Petrie, G. Soules, and N. Weiss. A maximization technique in the statistical analysis of probabilistic functions of Markov chains. *Ann. Math. Stat.*, 41:164–171, 1970.

[7] H.C.P. Berbee and R.C. Bradley. A limitation of Markov representation for stationary processes. *Stochastic Processes and their Applications*, 18:33–45, 1984.

[8] P.J. Bickel and Y. Ritov. *Inference in Hidden Markov Models I: Local Asymptotic Normality in the Stationary Case*. Technical Report 383, Department of Statistics, University of California, Berkeley, 1993.

[9] P. Billingsley. *Ergodic Theory and Information*. John Wiley & Sons, New York, 1964.

[10] D. Blackwell and L. Koopmans. On the identifiability problem for functions of finite Markov chains. *Ann. Math. Statist.*, 28:1011–1015, 1957.

[11] R.C. Bradley. An addendum to "A Limitation of Markov Representation for Stationary Processes". *Stochastic Processes and their Applications*, xx:xx, 19xx.

[12] R.W. Brockett. Stochastic realization theory and planck's law for black body radiation. *Ricerche di Automatica*, 10:344–362, 1979.

[13] P. Brodatz. *Texture: A Photographic Album for Artists and Designers*. Dover, New York, 1966.

[14] R. Chellappa and R.L. Kashyap. Texture synthesis using 2-D noncausal autoregressive models. *IEEE Trans. Acoustics, Speech, and Signal Processing*, 33:194–203, 1985.

[15] G.A. Churchill. Stochastic models in heterogeneous DNA sequences. *Bulletin of Mathematical Biology*, 51:79–94, 1989.

[16] F. Comets and B. Gidas. Parameter estimation for Gibbs distributions from partially observed data. *The Annals of Applied Probability*, 2:142–170, 1992.

[17] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm (with discussion). *J. Roy. Statist. Soc. (B)*, 39:1–38, 1977.

[18] H. Derin and H. Elliott. Modeling and segmentation of noisy and texturted images using Gibbs random fields. *IEEE Trans. Pattern Analysis and Machine Intellignece*, 9:39–55, 1987.

[19] S.W. Dharmadhikari. Sufficient conditions for a stationary process to be a function of a finite Markov chain. *Ann. Math. Statist.*, 34:1033–1041, 1963.

[20] I.M. Elfadel and R.W. Picard. *Gibbs random fields, co-occurrences, and texture modeling.* Technical Report, MIT Media Laboratory Perceptual Computing Group, 1993. to appear in IEEE Trans. Pattern Analysis and Machine Intelligence.

[21] J.M. Francos, A.Z. Meiri, and B. Porat. *A unified texture model based on a 2-D Wold like decomposition.* Technical Report, Department of Electrical Engineering, Technion-Israel Institute of Technology, Haifa 32000, Israel, 1992.

[22] D.R. Fredkin and J.A.Rice. Correlation functions of a function of a finite-state Markov process with application to channel kinetics. *Math. Biosciences*, 87:161–172, 1987.

[23] F.R. Fredkin and J.A.Rice. *Maximum Likelihood Estimation and Identification Directly from Single-Channel Recordings.* Technical Report 338, Department of Statistics, University of California, Berkeley, 1992.

[24] A. Frigessi and M. Piccioni. Parameter estimation for 2 dimensional Ising fields corrupted by noise. *Stochastic Processes and their Applications*, 34:297–311, 1990.

[25] A. Gagalowicz and S.D. Ma. Sequential synthesis of natural textures. *Computer Vision, Graphics, and Image Processing*, 30:289–315, 1985.

[26] S. Geman, A. Kehagias, and H. Künsch. *Consistent Estimation of Stationary Processes and Stationary Random Fields.* Technical Report, Division of Applied Mathematics, Brown University, Providence, RI 02912, 1993.

[27] H.-O. Georgii. *Gibbs Measures and Phase Transitions.* Walter de Gruyter, New York, 1988.

[28] E.J. Gilbert. On the identifiability problem for functions of finite Markov chains. *Ann. Math. Statist.*, 30:688–697, 1959.

[29] U. Grenander. *Abstract Inference.* John Wiley & Sons, New York, 1978.

[30] H. Ito, S.-I. Amari, and K. Kobayashi. *Identifiability of Hidden Markov Information Sources and Their Minimum Degrees of Freedom.* Technical Report METR 90-07, Department of Mathematical Engineering and Instrumentation Physics, University of Tokyo, 1990.

[31] C. Ji. *Sieve Estimators for Pair-Interaction Potentials and Local Characteristics in Gibbs Random Fields.* Technical Report 2037, Institute of Statistics, University of North Carolina, 1990.

[32] C. Ji and L. Seymour. *On the Selection of Markov Random Field Texture Models.* Technical Report, Department of Statistics, University of North Carolina, 1992.

[33] G. Kallianpur. *Stochastic Filtering Theory.* Springer-Verlag, New York, 1980.

[34] A. Kehagias. *Approximation of Stochastic Processes by Hidden Markov Models.* PhD thesis, Division of Applied Mathematics, Brown University, 1992.

[35] P. Kienker. Equivalence of aggregated Markov models of ion-channel gating. *Proc. R. Soc. Lond. B*, 236:269–309, 1989.

[36] R. Kindermann and J.L. Snell. *Markov Random Fields and their Applications.* American Mathematical Society, Providence, RI, 1980.

[37] A. Krogh, M. Brown, I.S. Mian, K. Sjölander, and D. Haussler. *Hidden Markov models in computational biology: applications to protein modeling.* Technical Report, Computer and Information Sciences, University of California, Santa Cruz, 1992.

[38] D.S. Ornstein and B. Weiss. How sampling reveals a process. *Ann. Probab.*, 18:905–930, 1990.

[39] S.M. Pincus. Approximating Markov chains. *Proc. Natl. Acad. Sci.*, 89:xxx, 1992.

[40] W. Qian and D.M. Titterington. Estimation of parameters in hidden Markov models. *Phil. Trans. R. Soc. Lond. A*, 337:407–428, 1991.

[41] L.R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77:257–286, 1989.

[42] M. Rosenblatt. *Markov Processes. Structure and Asymptotic Behavior.* Springer-Verlag, New York, 1971.

[43] L. Younes. Parametric inference for imperfectly observed Gibbsian fields. *Probability Theory and Related Fields*, xx:xxx–xxx, 19xx.

Figure Legends.

Figure 1. Panel A: xxx millisecond segment from the phoneme [ah]. Panel B: xxx millisecond segment from the phoneme [nn].

Figure 2. Panel A: HMM estimated from data in Figure 1A. Panel B: HMM estimated from data in Figure 1B.

Figure 3. Panel A: thresholded image of straw. Panel B: thresholded image of paper.

Figure 4. Panel A: HMM estimated from data in Figure 3A. Panel B: HMM estimated from data in Figure 3B.

A

B

Figure 1

A

B

Figure 2

A

B

Figure 3

A

B

Figure 4

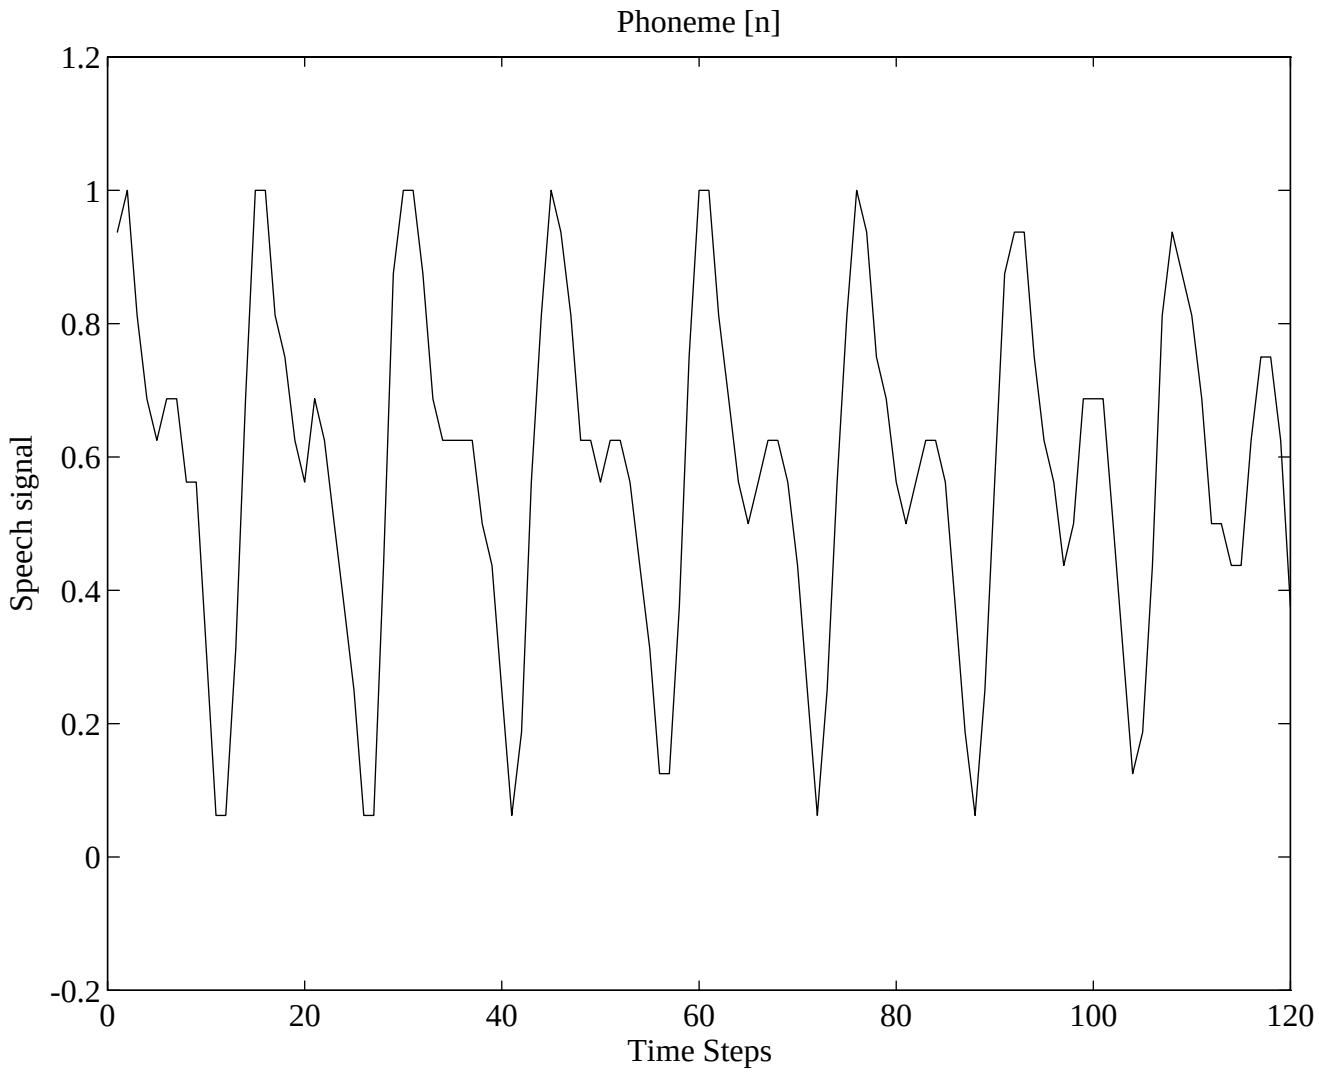
Phoneme [n]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

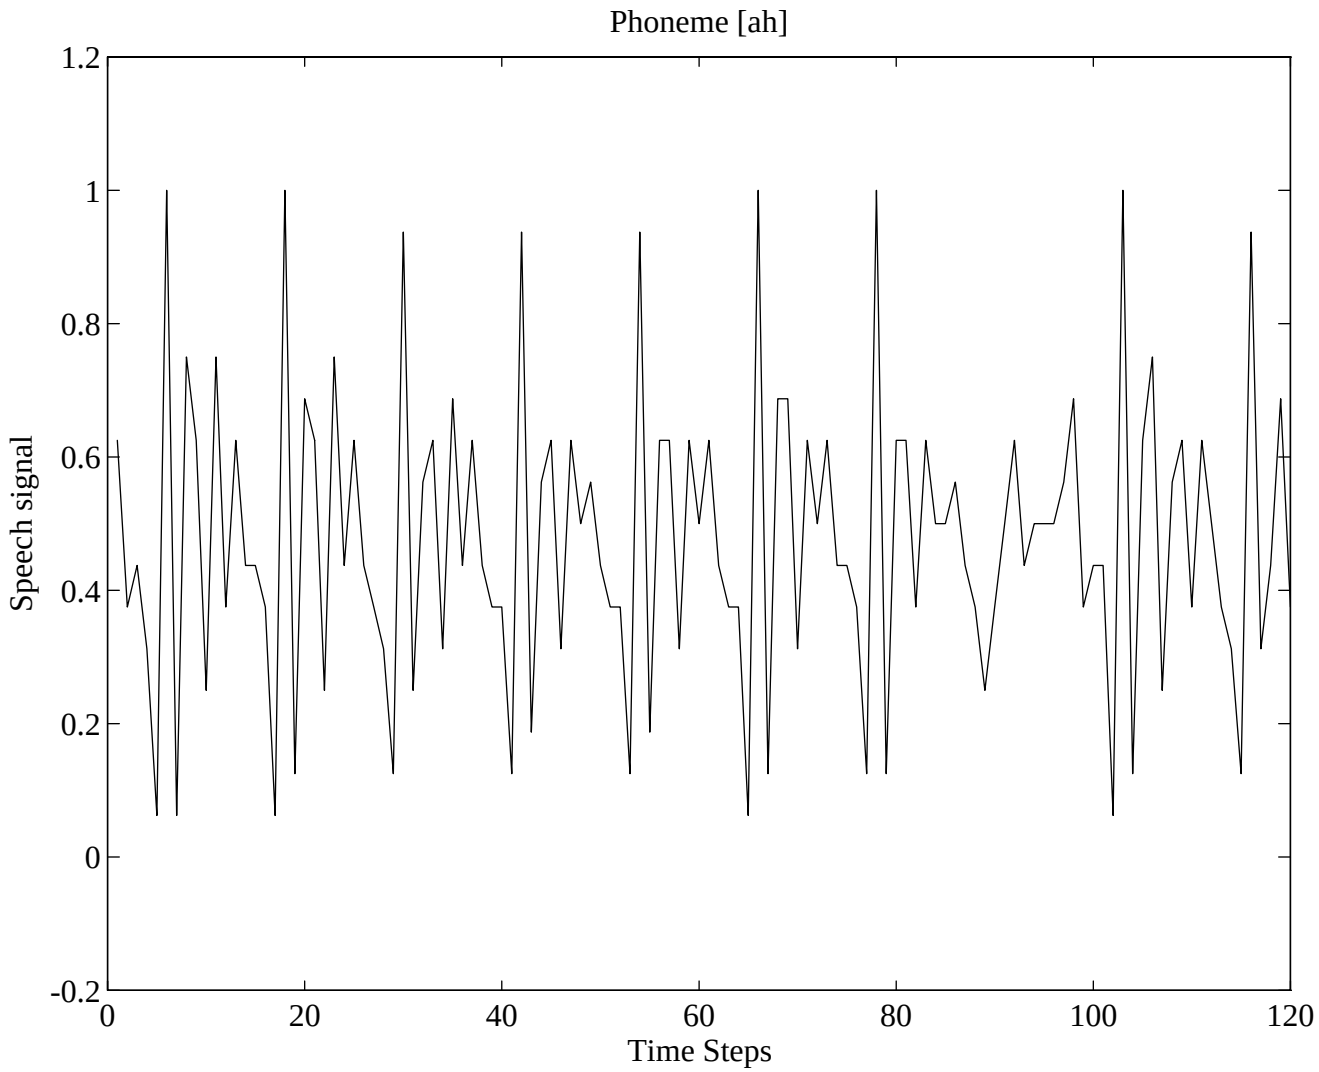
Phoneme [ah]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

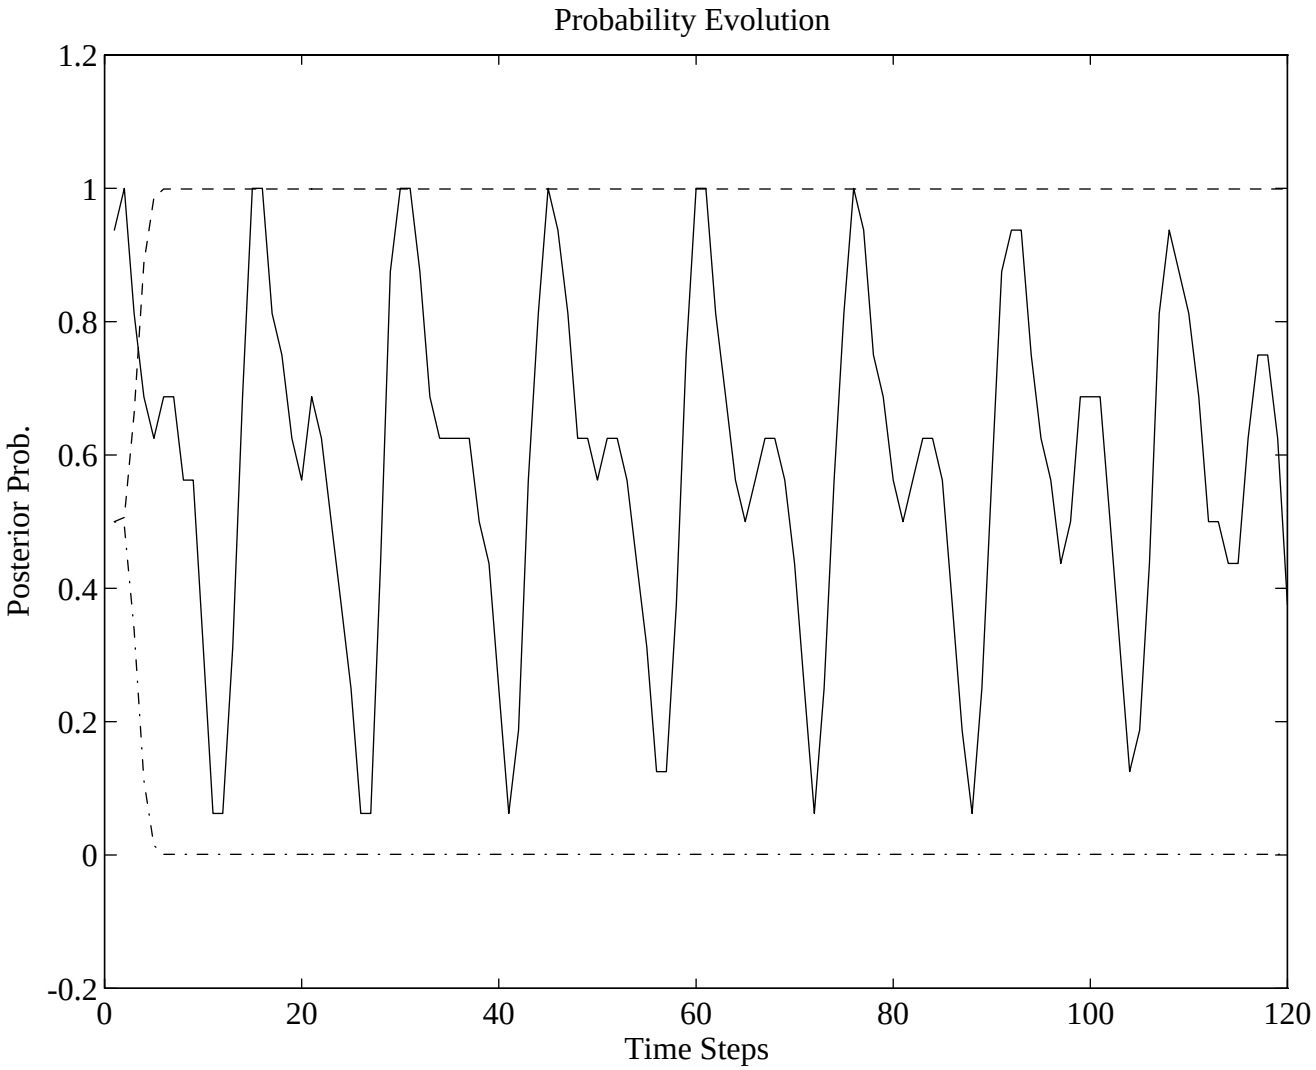
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

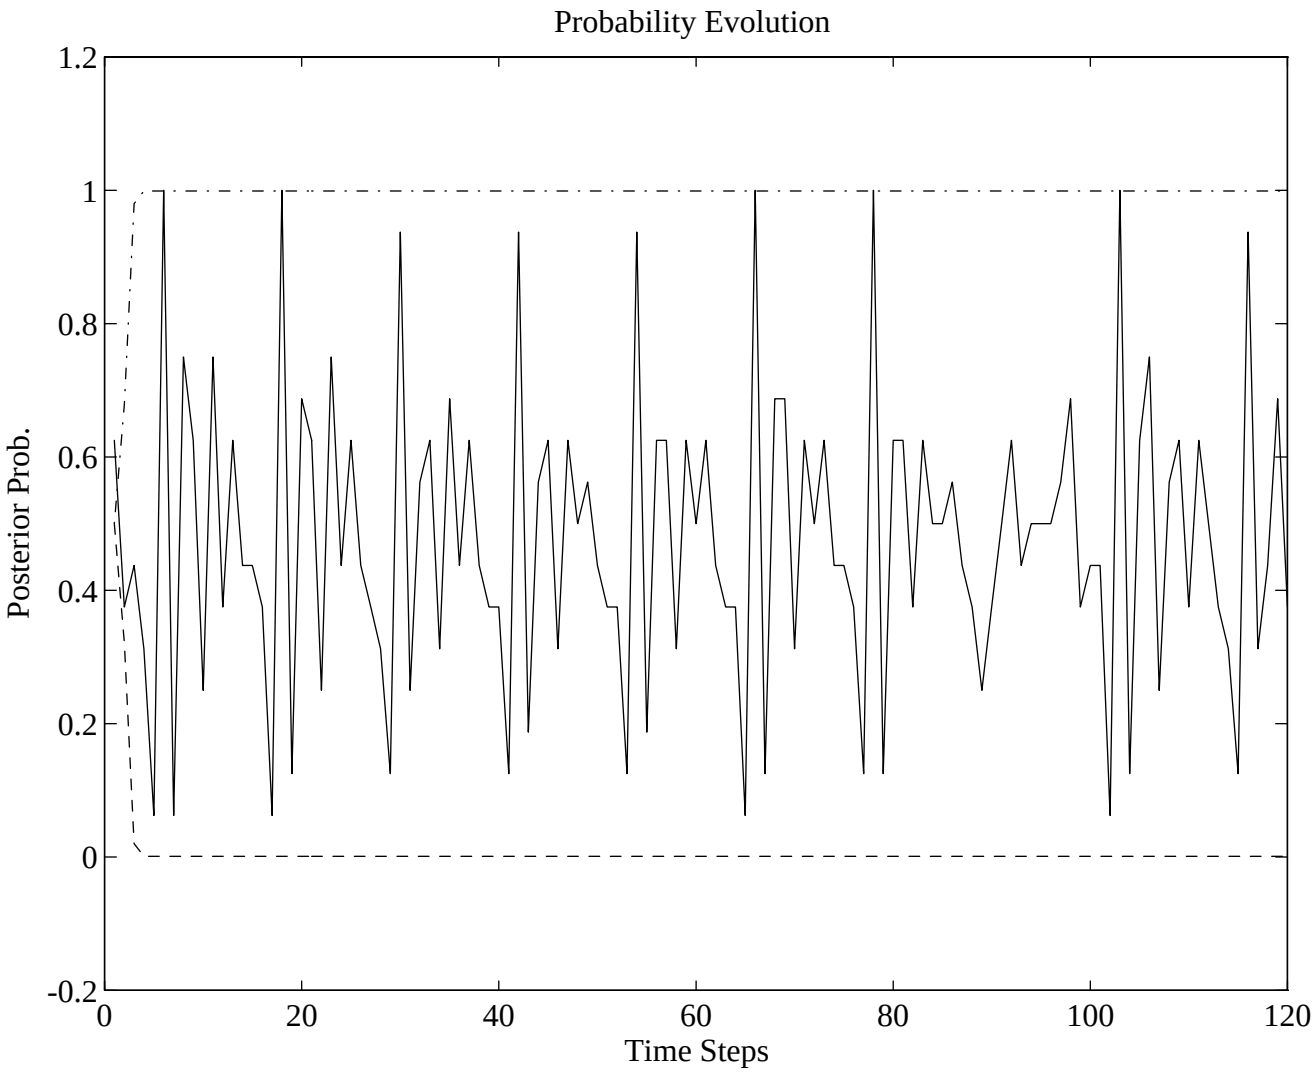
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

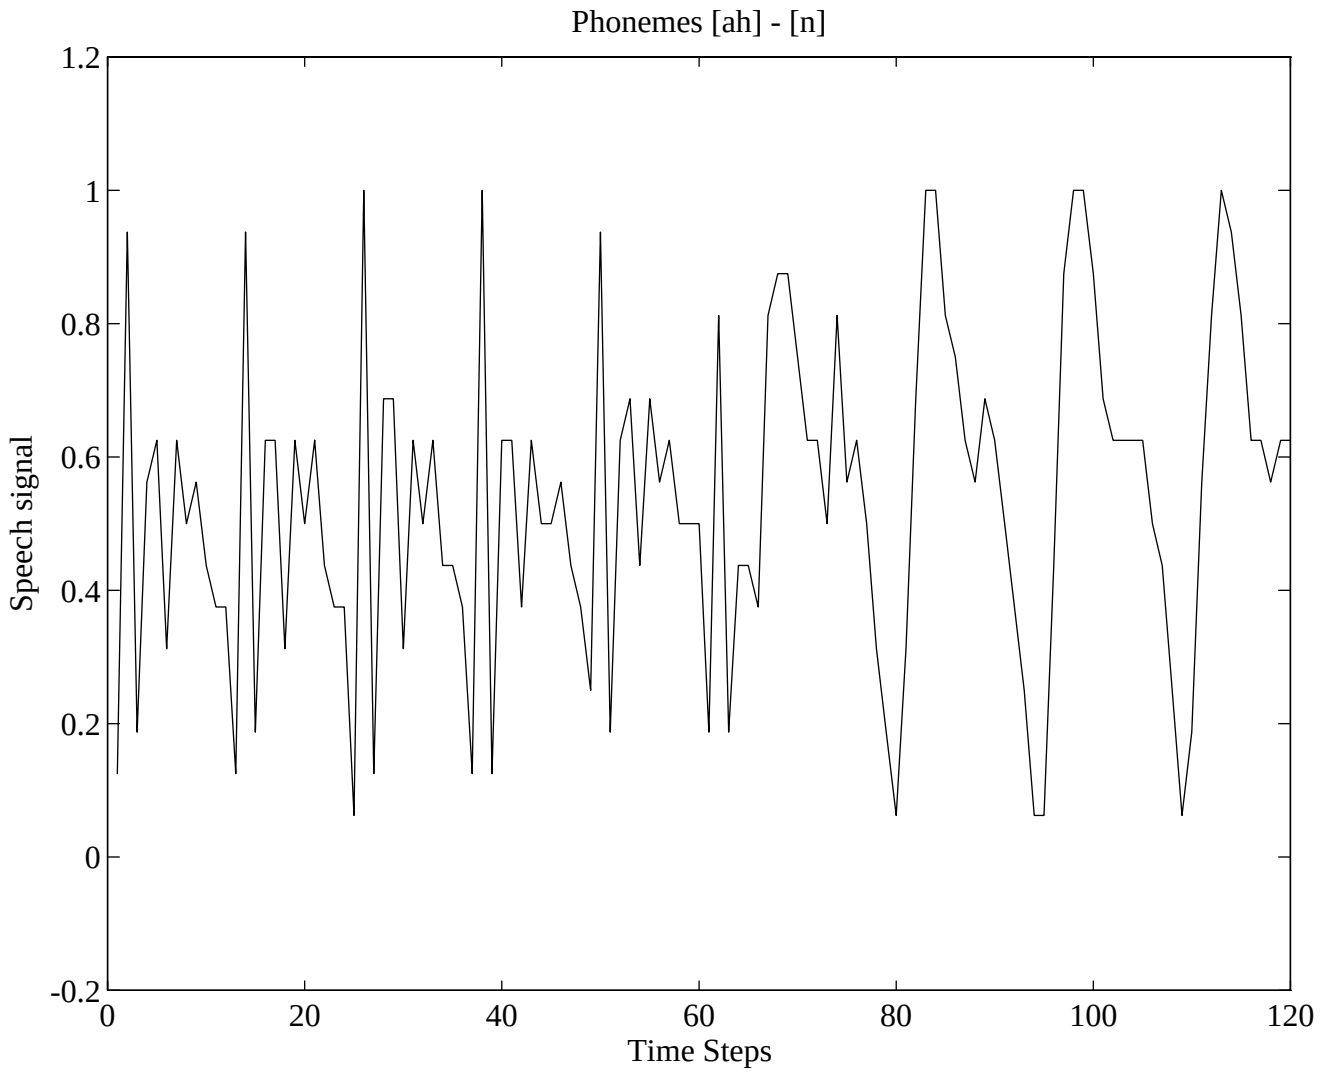
Phonemes [ah] - [n]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

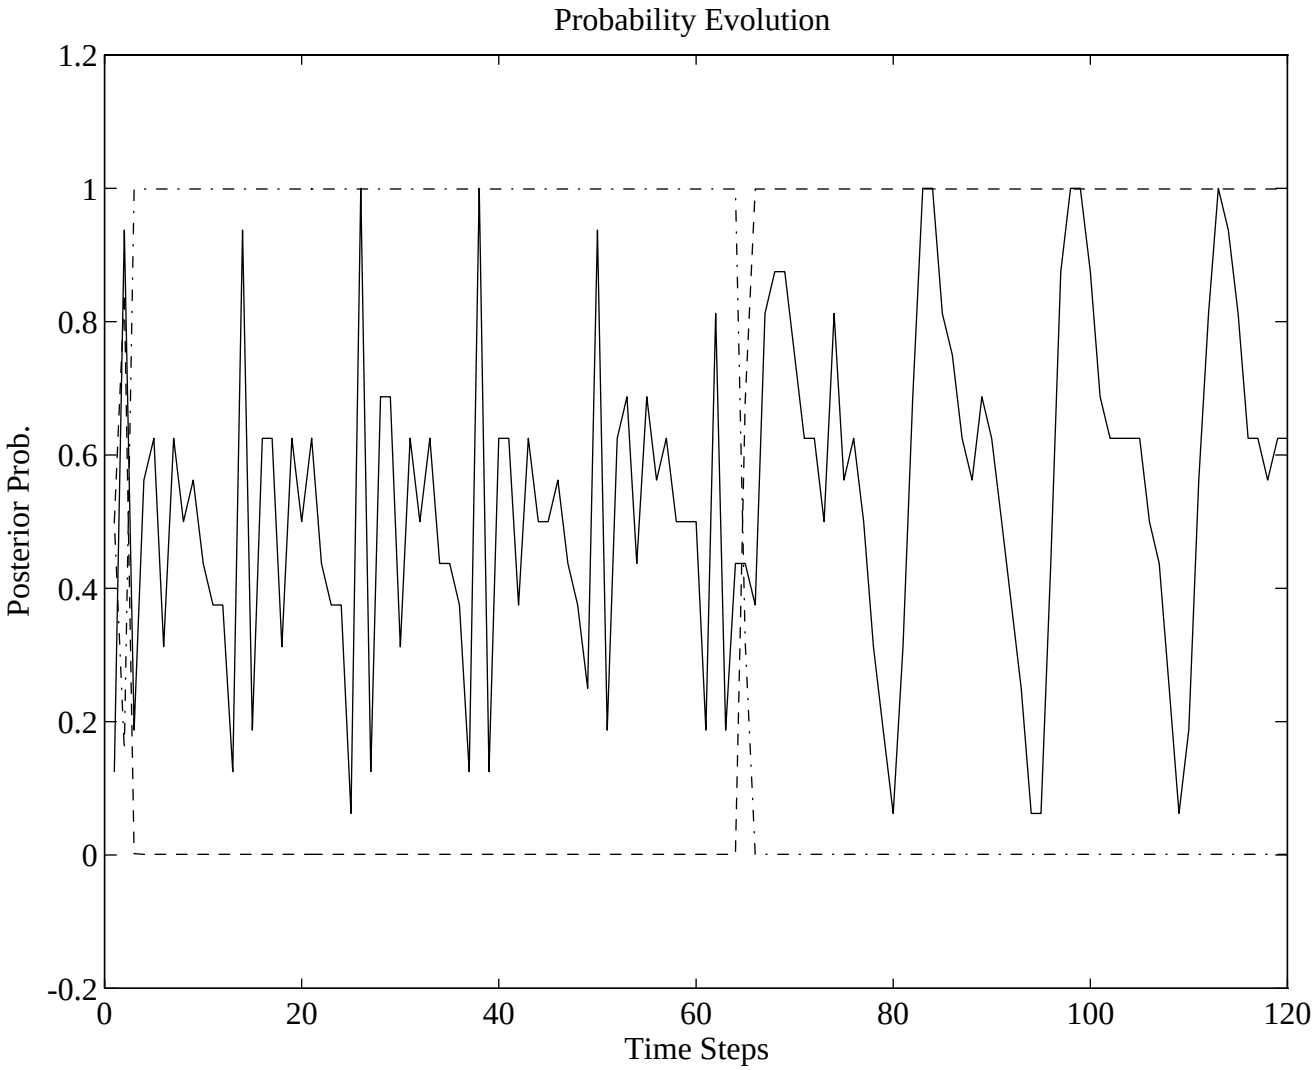
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

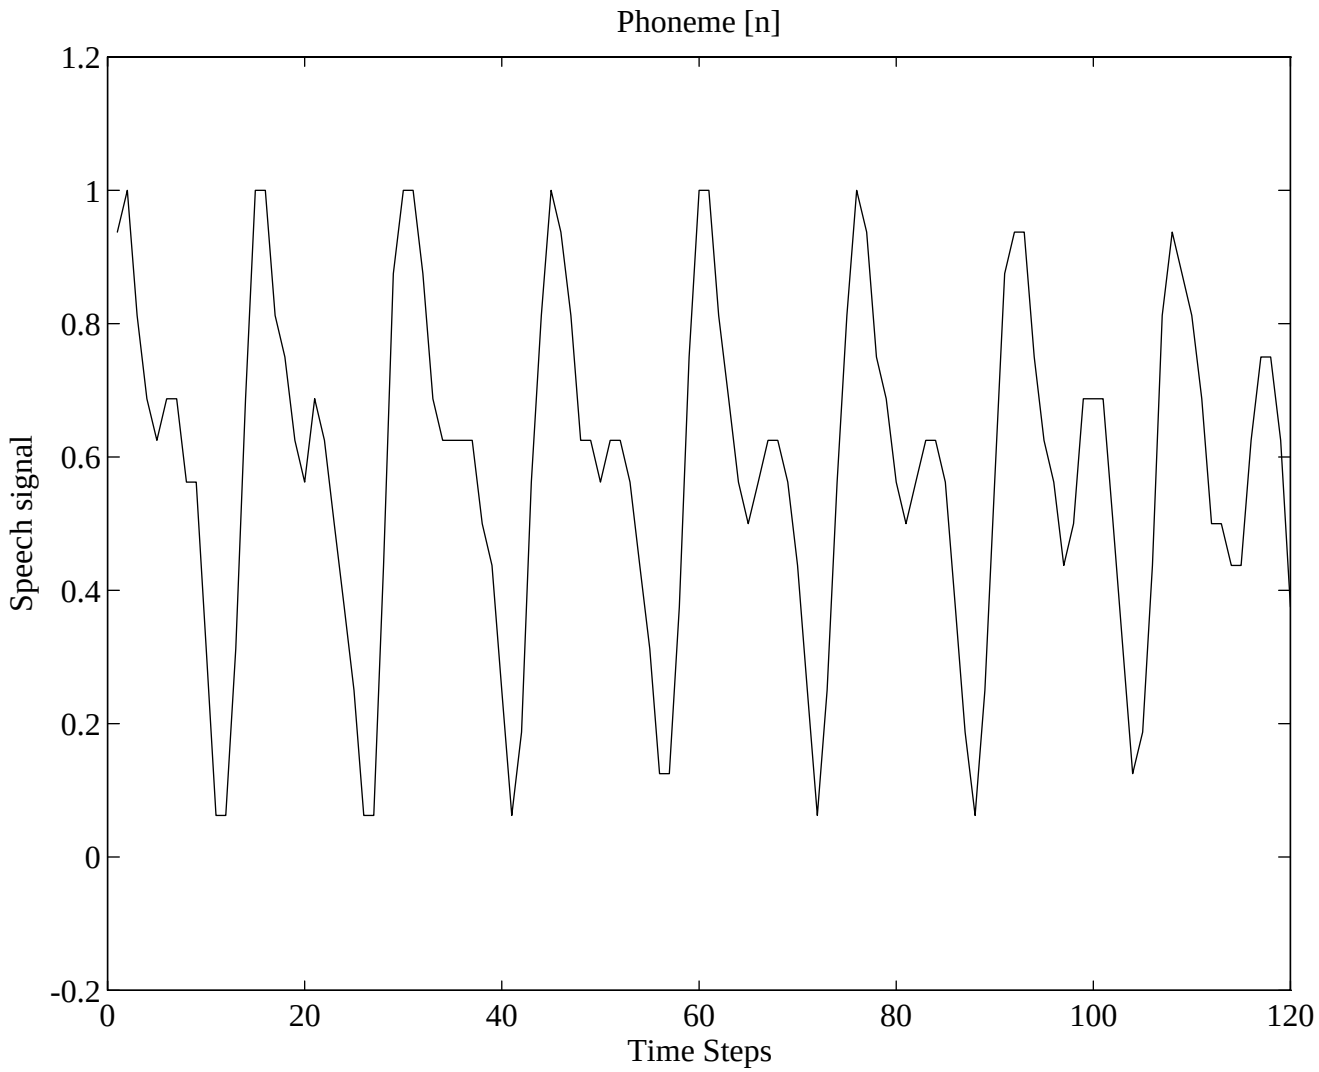
Phoneme [n]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

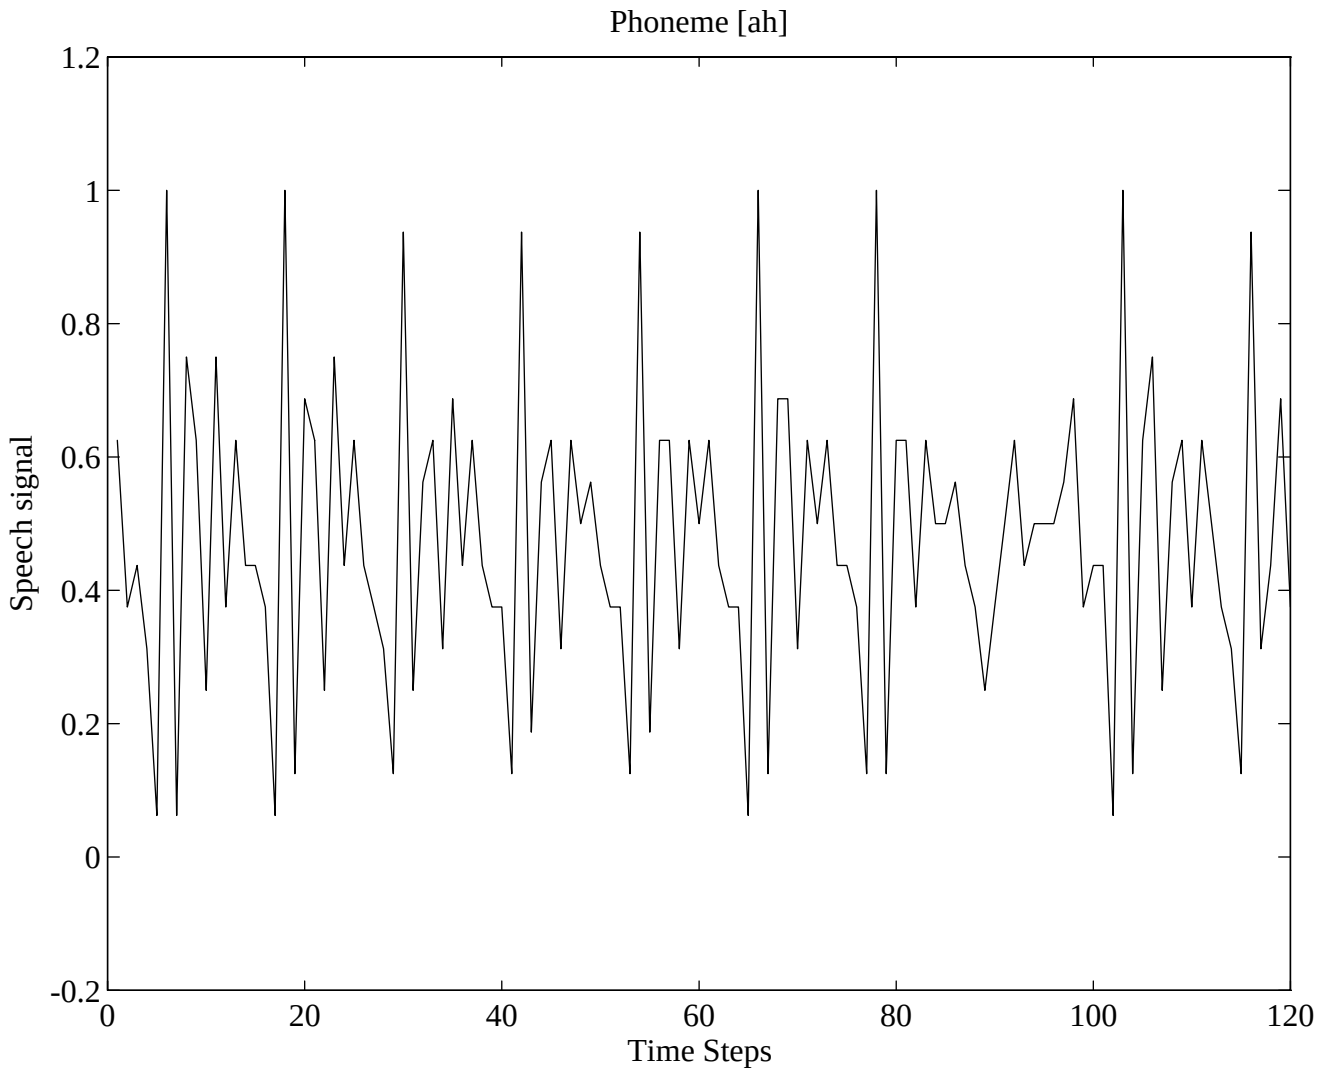
Phoneme [ah]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

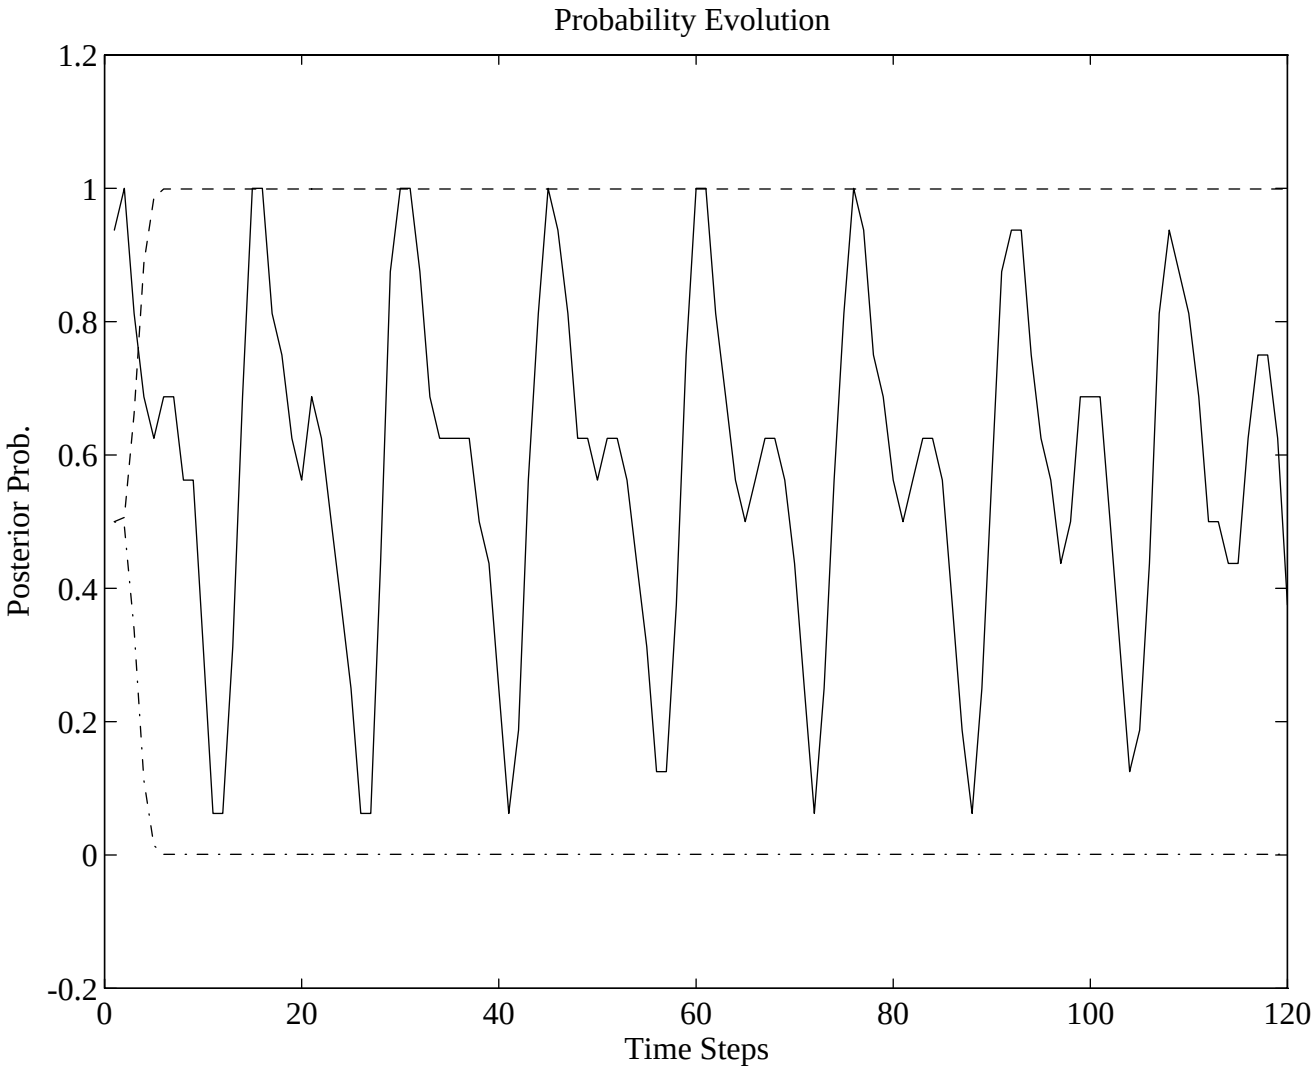
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

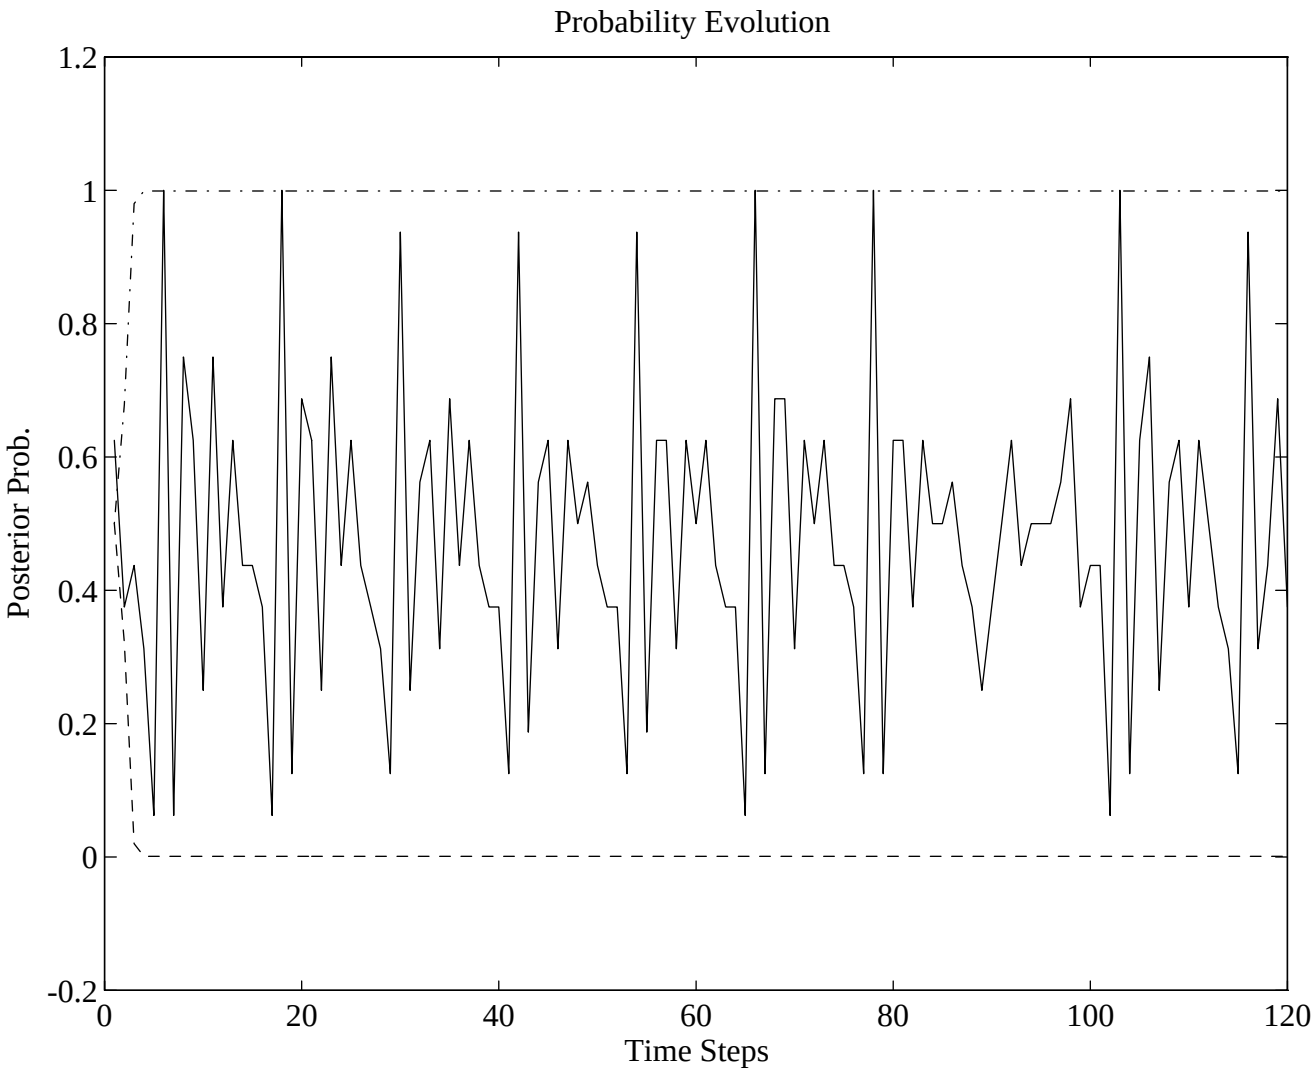
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

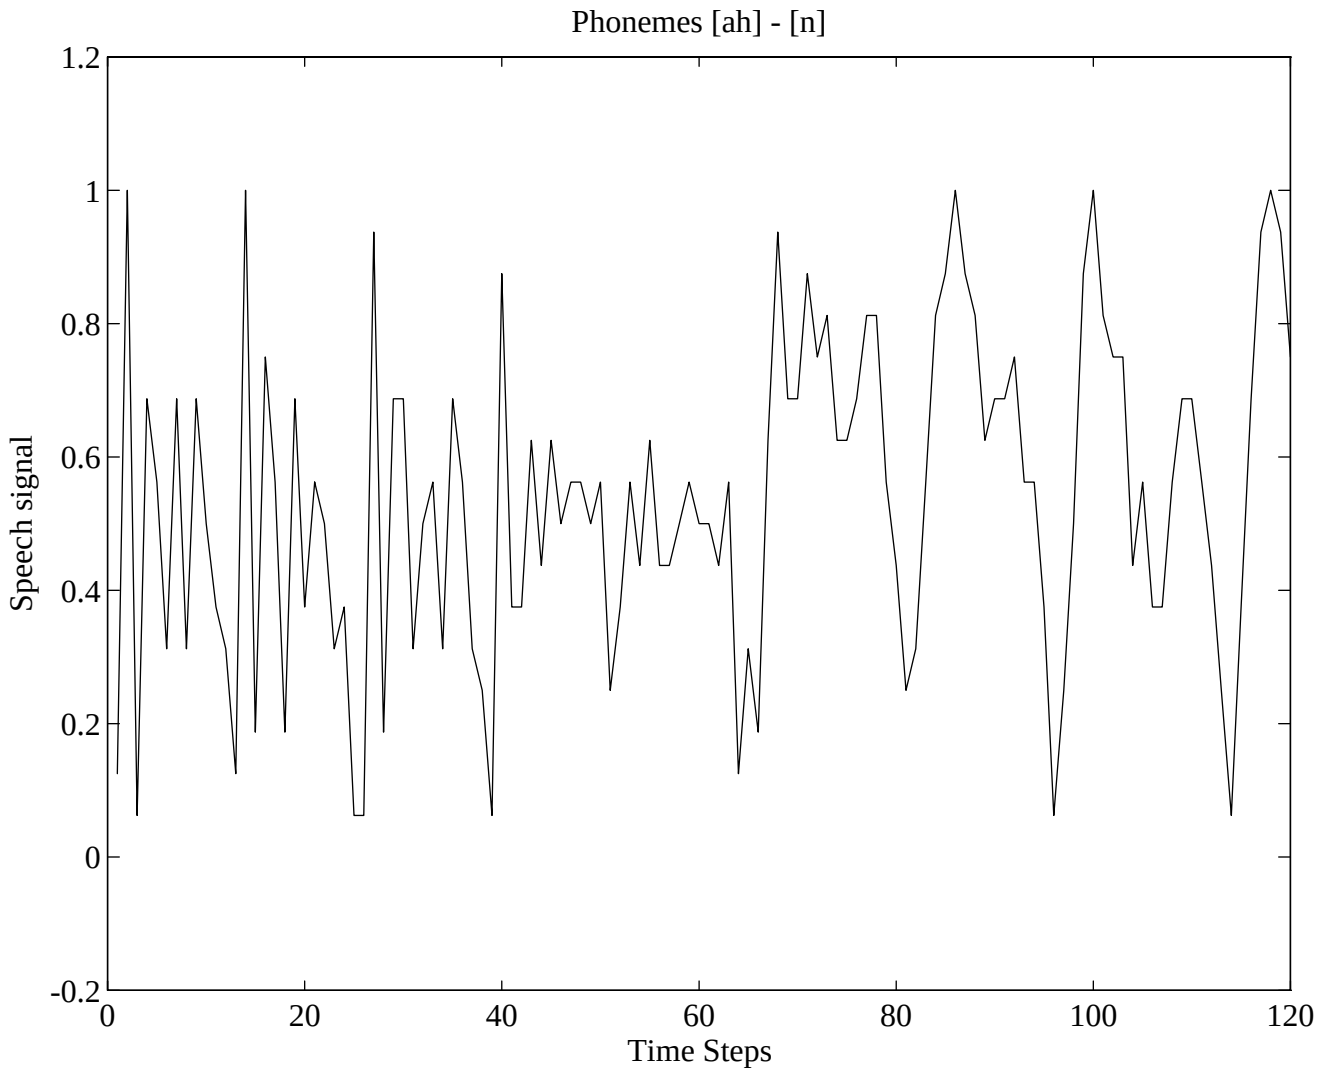
Phonemes [ah] - [n]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

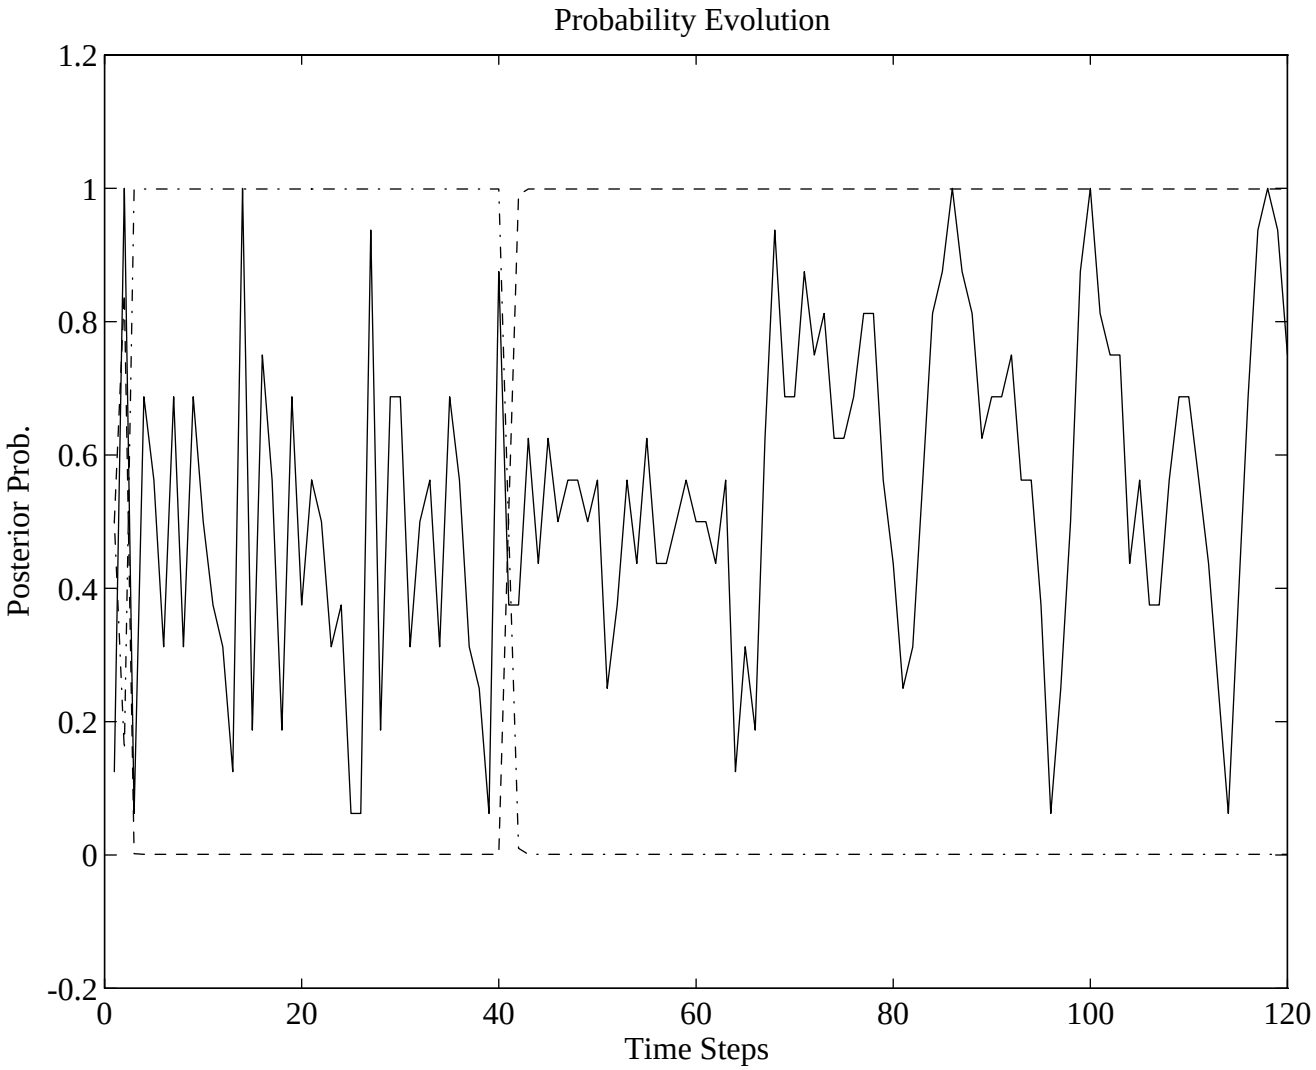
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

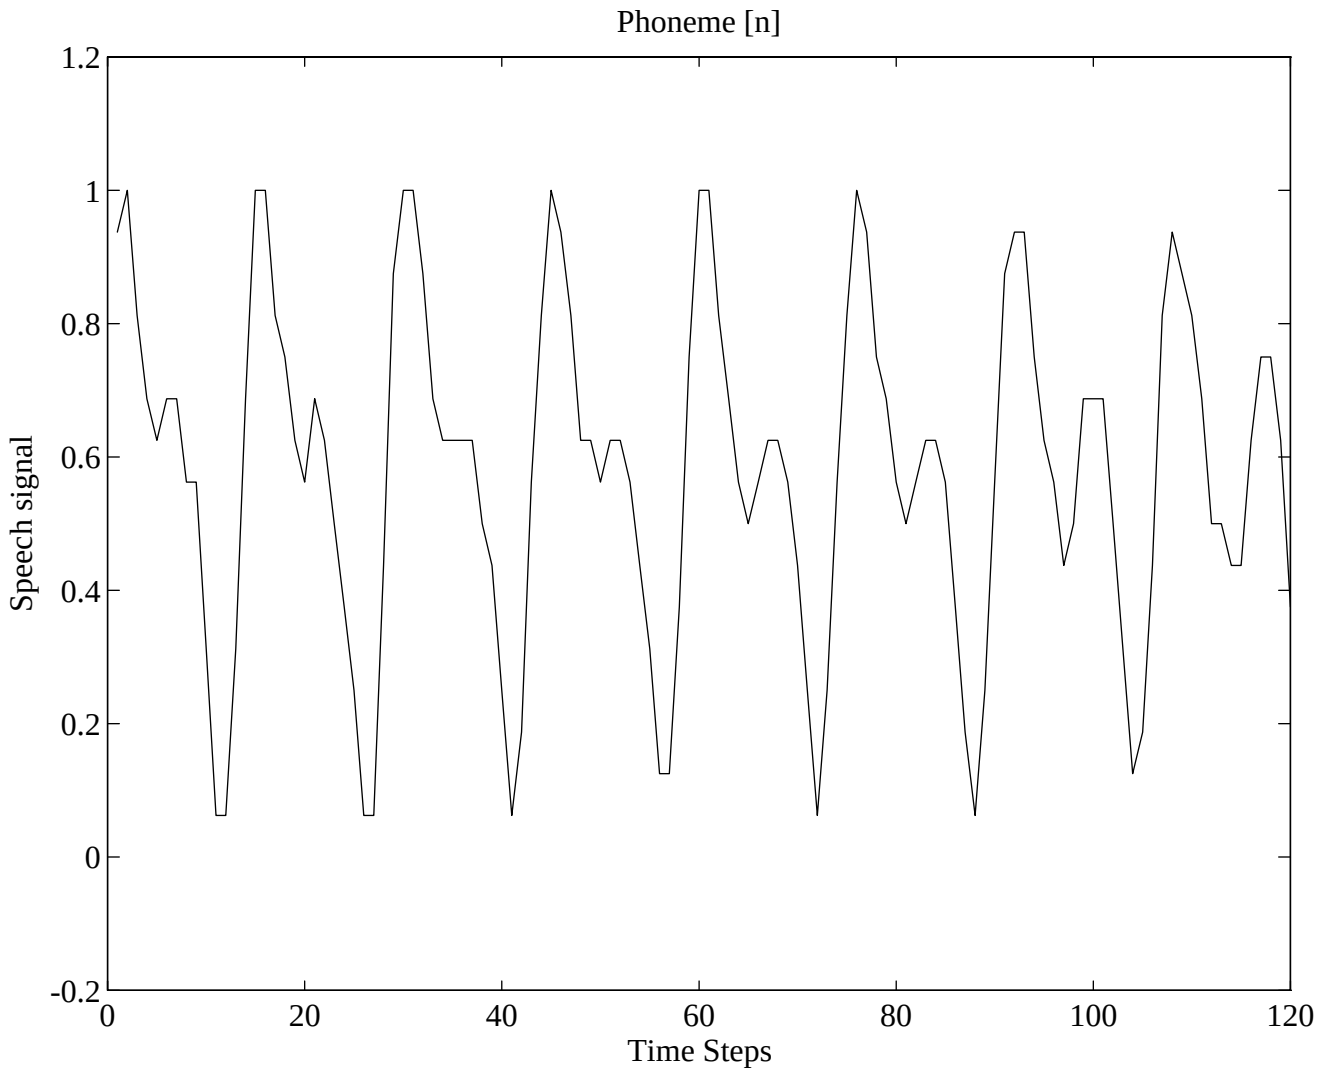
Phoneme [n]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

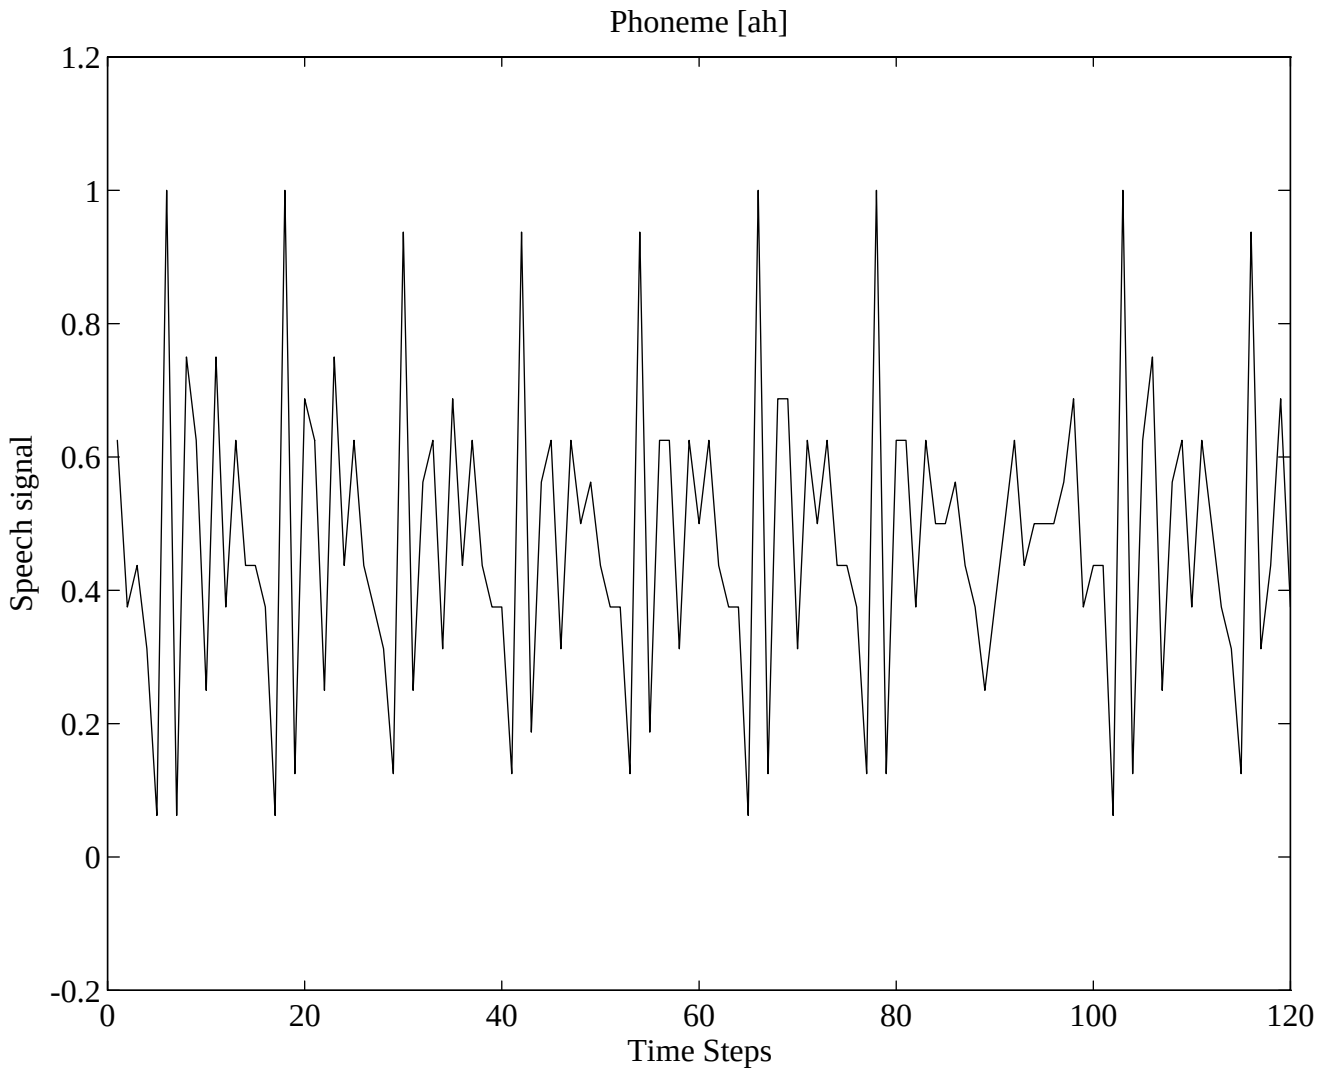
Phoneme [ah]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

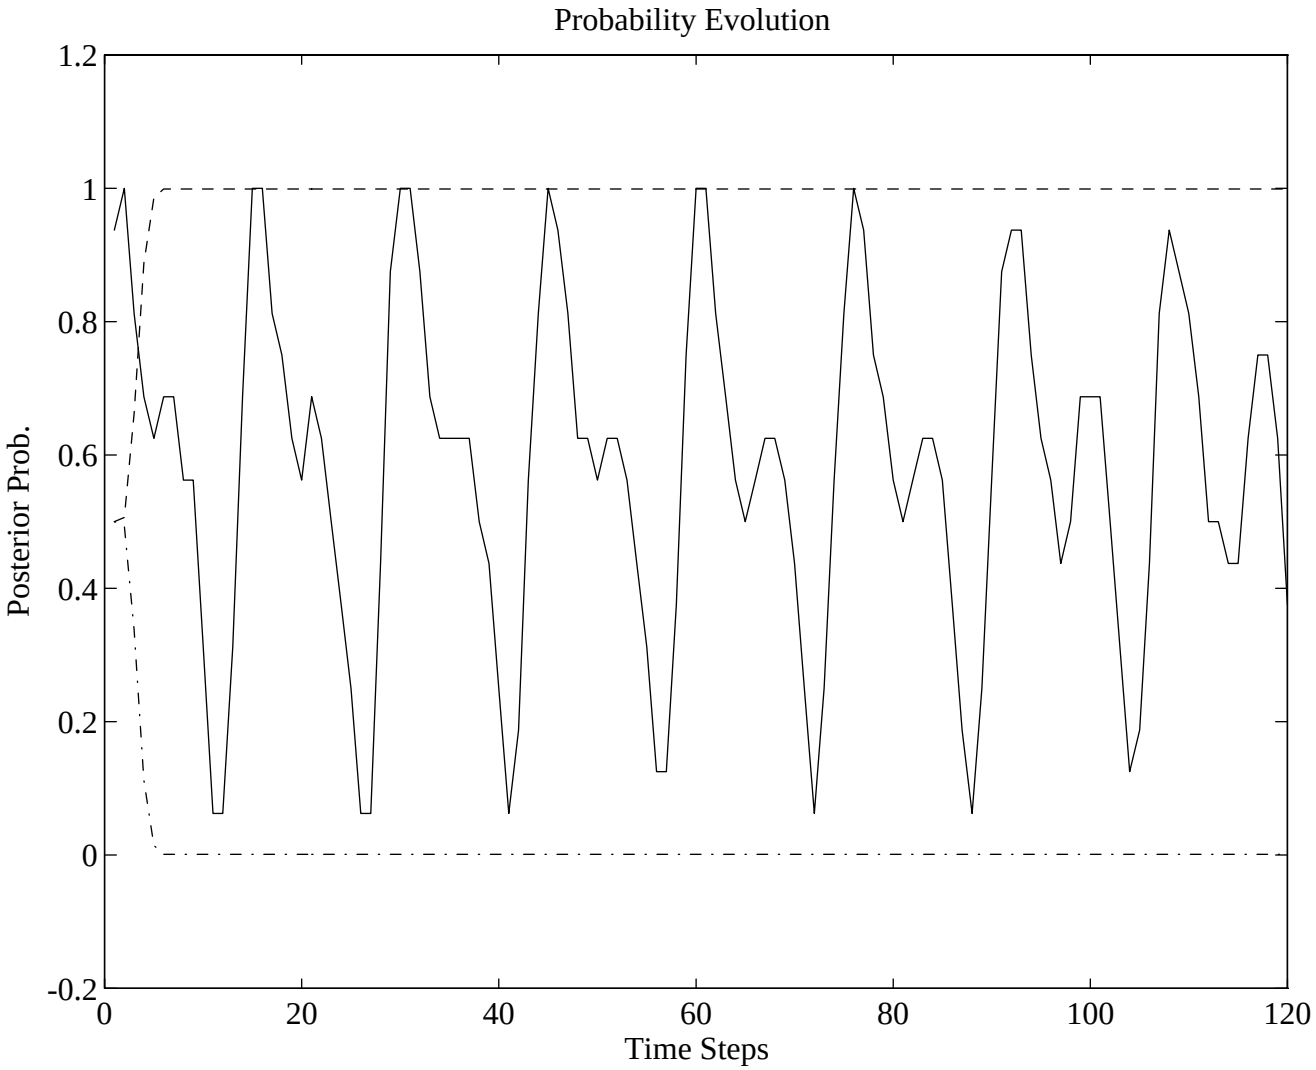
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

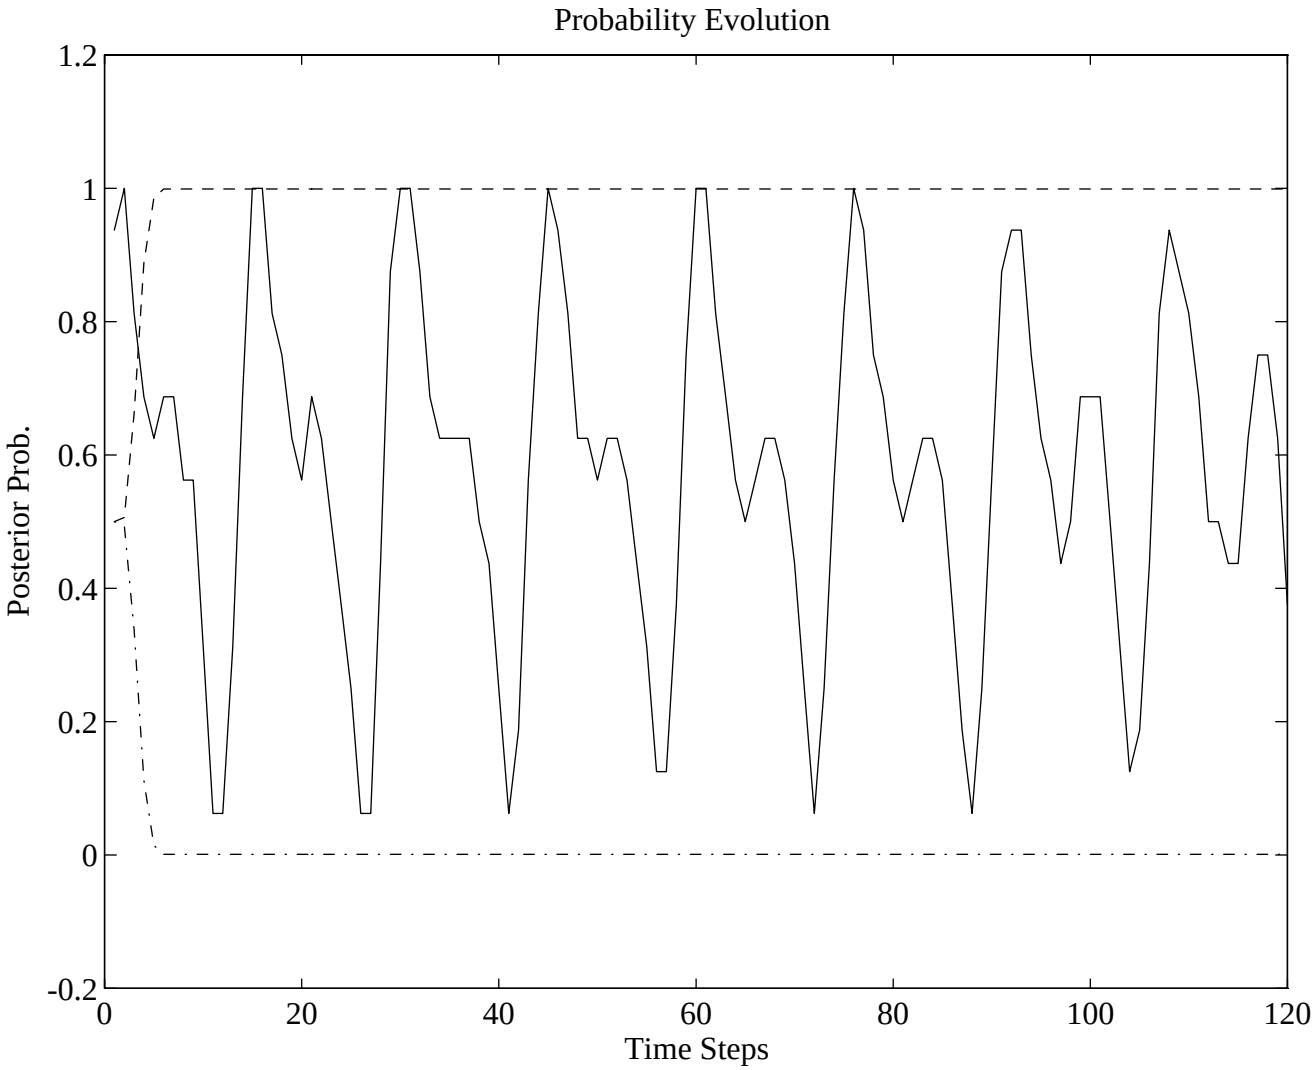
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

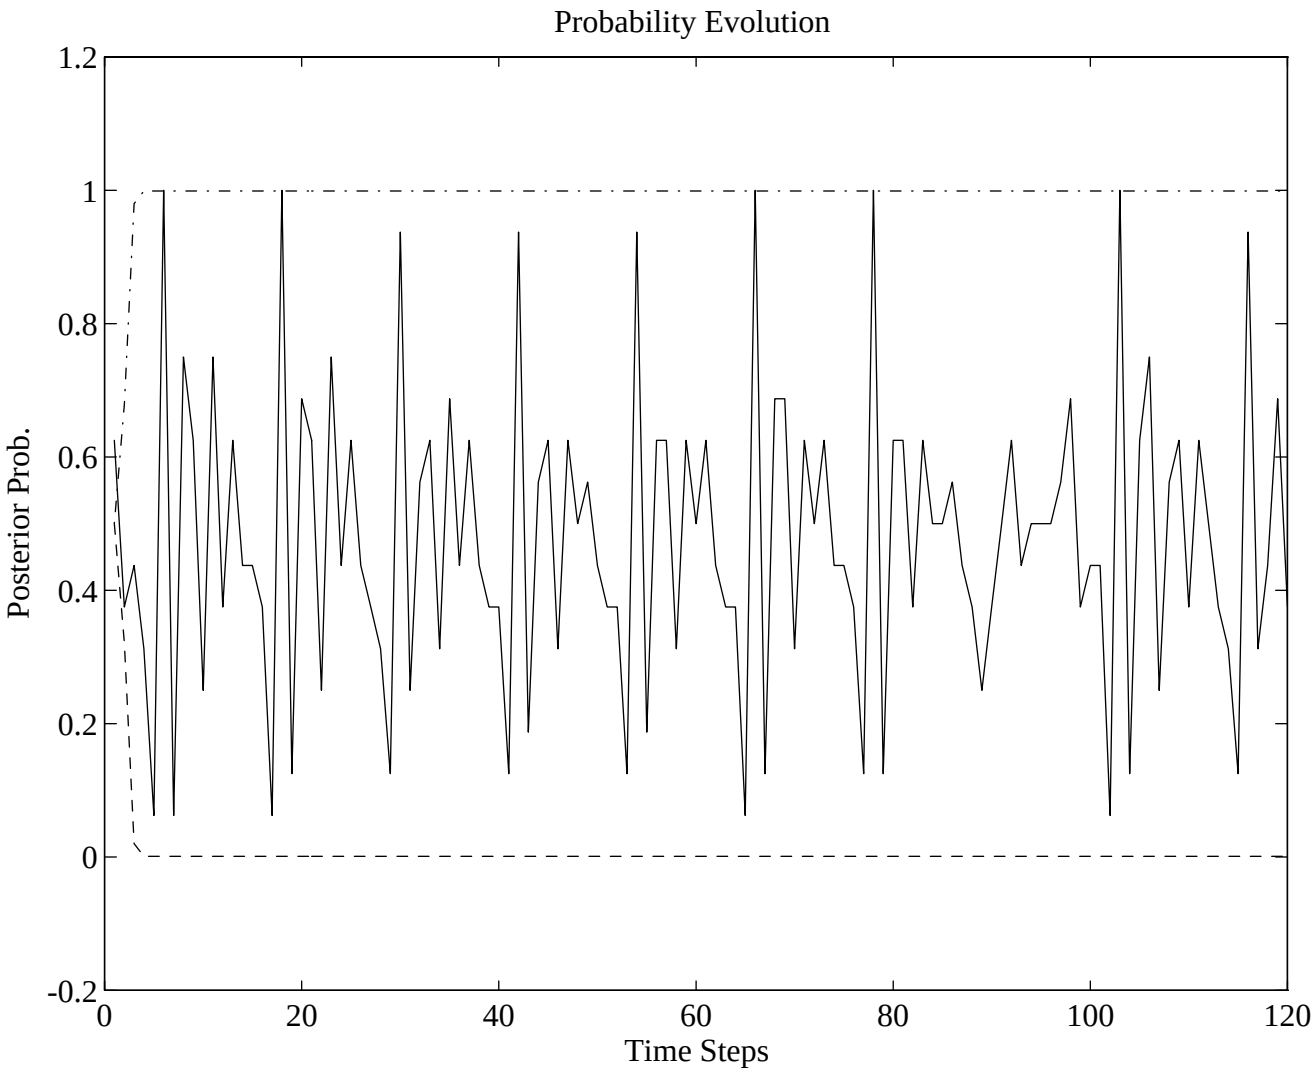
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

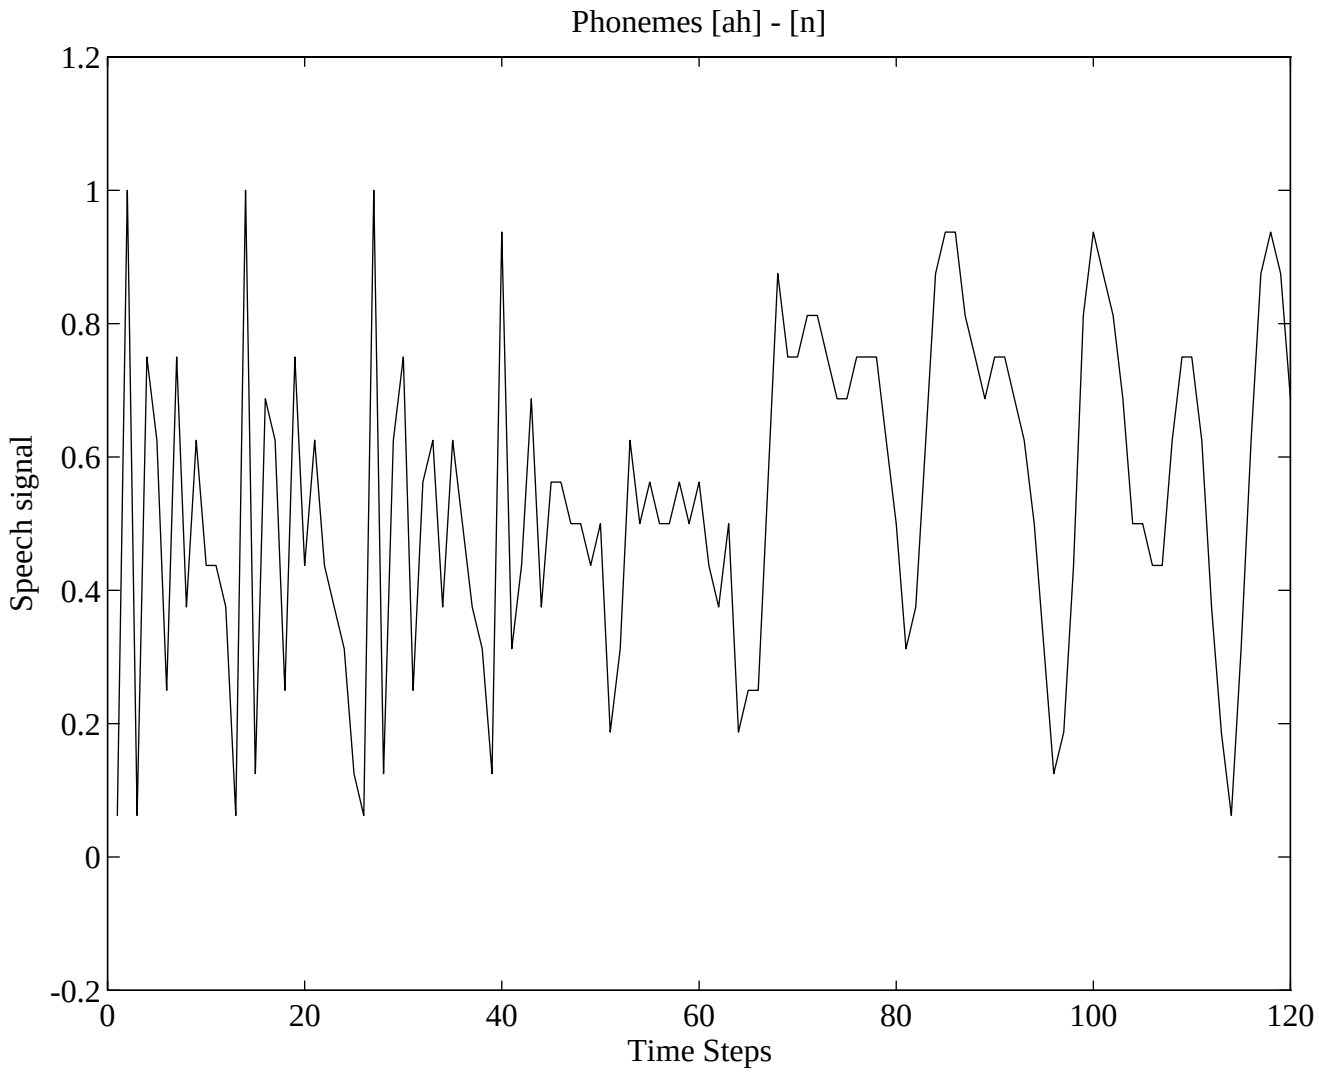
Phonemes [ah] - [n]
Speech signal
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120

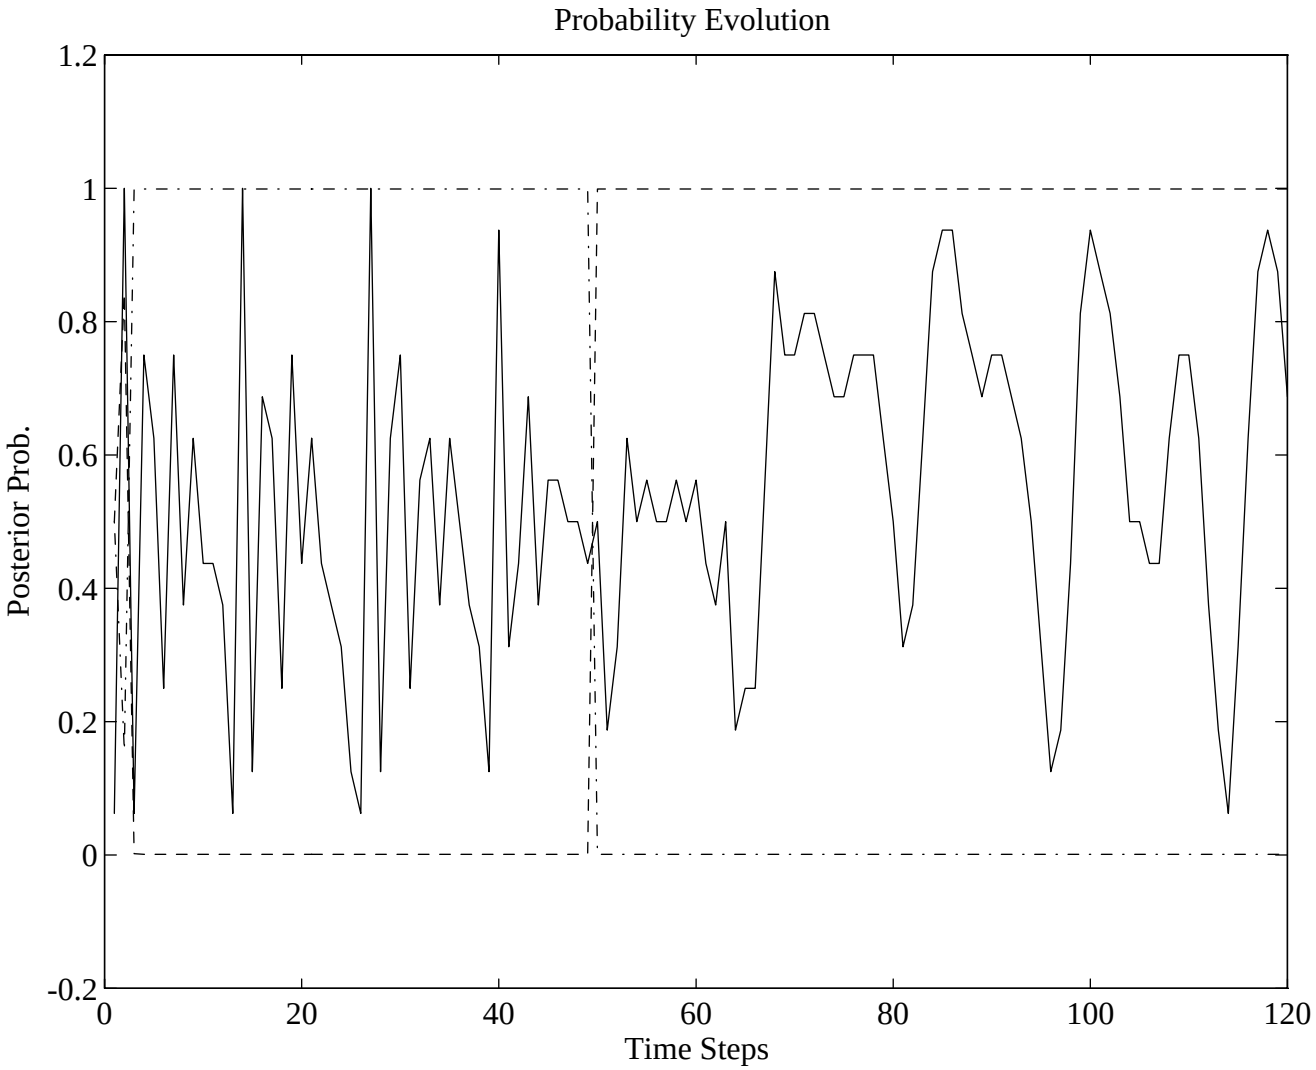
Probability Evolution
Posterior Prob.
Time Steps
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
0
20
40
60
80
100
120