Available online at www.sciencedirect.com
SCIENCE DIRECT®
Information Processing Letters 95 (2005) 521–528
Information Processing Letters
www.elsevier.com/locate/ipl

Isolated word recognition with the *Liquid State Machine*: a case study

D. Verstraeten *, B. Schrauwen, D. Stroobandt, J. Van Campenhout

Department of Electronics and Information Systems, Ghent University, Sint-Pietersnieuwstraat 41, 9000 Gent, Belgium

Available online 14 July 2005

Abstract

The Liquid State Machine (LSM) is a recently developed computational model with interesting properties. It can be used for pattern classification, function approximation and other complex tasks. Contrary to most common computational models, the LSM does not require information to be stored in some stable state of the system: the inherent dynamics of the system are used by a memoryless readout function to compute the output. In this paper we present a case study of the performance of the Liquid State Machine based on a recurrent spiking neural network by applying it to a well known and well studied problem: speech recognition of isolated digits. We evaluate different ways of coding the speech into spike trains. In its optimal configuration, the performance of the LSM approximates that of a state-of-the-art recognition system. Another interesting conclusion is the fact that the biologically most realistic encoding performs far better than more conventional methods.

Keywords: Liquid State Machine; Speech recognition; Spiking neural networks; Parallel processing

1. Introduction

Many complex computational problems have a strong temporal aspect: not only the value of the inputs is important, but also their specific sequence and precise occurrence in time. Tasks such as speech recognition, object tracking, robot control or biometrics are inherently temporal, as are many of the tasks that are usually viewed as 'requiring intelligence'. However, most computational models do not explicitly take the temporal aspect of the input into account or transform the time-dependent inputs to static input using, e.g., a tapped delay line. These methods disregard the temporal information contained in the inputs in two ways: the time-dependence of the inputs within a certain time window is compressed into a static *snapshot* and is therefore partially lost, and the temporal correlation between different windows is not preserved either.

The Liquid State Machine (LSM) [1] avoids these problems by construction. The LSM is a computational concept (its structure is depicted in Fig. 1): a reservoir of recurrently interacting nodes is stim-

* Corresponding author.
E-mail addresses: David.Verstraeten@ugent.be (D. Verstraeten), Benjamin.Schrauwen@elis.ugent.be (B. Schrauwen), Dirk.Stroobandt@ugent.be (D. Stroobandt), Jan.VanCampenhout@ugent.be (J. Van Campenhout).

0020-0190/$ – see front matter
doi:10.1016/j.ipl.2005.05.019