Modeling consciousness

Frédéric Dandurand and Thomas R. Shultz
Department of Psychology, McGill University, Montréal, Québec, H3A 1B1 Canada. **F.Dandurand@sympatico.ca** **thomas.shultz@mcgill.ca**
http://www.psych.mcgill.ca/perpg/fac/shultz/default.html

Abstract: Perruchet & Vinter do not fully resolve issues about the role of consciousness and the unconscious in cognition and learning, and it is doubtful that consciousness has been computationally implemented. The cascade-correlation (CC) connectionist model develops high-order feature detectors as it learns a problem. We describe an extension, knowledge-based cascade-correlation (KBCC), that uses knowledge to learn in a hierarchical fashion.

Issues concerning the role of consciousness in cognition and learning have bedeviled psychology for many years. It has proved to be enormously difficult to delineate the role of consciousness in various psychological phenomena. Perruchet & Vinter (P&V) show that a theory emphasizing the unconscious processing of information about which people are conscious accounts for a wide range of phenomena just as well as a theory that allows for unconscious processing of information about which people are not conscious. Or, put more simply, assuming conscious awareness of information fits the data just as well as assuming an unconscious does. Although this is useful and even interesting, it does little to resolve the classical issues of the roles of consciousness and the unconscious in cognition and learning. The idea that multiple theories account for a range of phenomena is common in psychology and remains true in this domain.

With their PARSER simulations, P&V bring a potentially useful tool to the arena. The critical importance of attentional focus in this model is extremely interesting. As P&V correctly note, however, computational models are essentially neutral with respect to consciousness because none of them actually have it. An important reason why none of them have it is that no one yet knows enough about consciousness to implement it computationally. Equating consciousness with attention, as in PARSER, is an interesting gambit, but it is far from clear that consciousness has been fully and effectively implemented in PARSER. Many researchers would insist on some level of awareness as a key marker for consciousness.

We note with interest that P&V intend to implement PARSER in a connectionist framework, specifically with CC algorithm (Fahlman & Lebiere 1990), which we have used in a wide variety of psychological simulations (Shultz, in press). A recent simulation concerned a data set given extensive coverage by P&V, that of Marcus et al. (1999). This simulation showed that a cascade-correlation encoder network covers the major features of Marcus et al.'s infant data (Shultz & Bale 2001) and generalizes outside the range of the training patterns, at the same time showing psychologically plausible content effects. Like most connectionist models, this one requires neither conscious awareness nor explicitly built-in attentional focus. Attention is not a cause of the model's successful learning but a result of learning – the networks learn what features of the stimuli to attend to.

Perruchet and Vinter foresee two problems in implementing PARSER using CC: (1) the representations embedded in the connection weights between units are not formatted to serve as new coding primitives; and (2) it is difficult to implement the idea that associations apply to increasingly complex representations. Despite the fact that CC develops high-order feature detectors as it learns a problem, P&V correctly point out that explicit associations between simple and complex representations are not modeled using standard CC. Our extension of CC (KBCC) explicitly addresses the problem of knowledge re-use and the building of complex representations (Rivest & Shultz 2002; Shultz & Rivest 2001; Thivierge & Shultz 2002).

Both KBCC and CC are constructive methods in which the network topology expands as necessary for learning the problem at hand. Expansion occurs when a new unit from a pool of candidates is added into the network. The main improvement of KBCC over CC is in the content of the pool of candidates used for recruitment. Whereas CC uses only simple units, KBCC extends the pool of candidates to arbitrarily complex CC networks. Knowledge contained in previously trained CC networks is therefore directly available for use in new problems, hence implementing knowledge re-use and solving the first issue raised by P&V.

Although KBCC can recruit any kind of sub-network if it correlates best with the residual network error, there are three particularly useful sources of knowledge: sub-tasks (e.g., rules for DNA splicing; Thivierge & Shultz 2002), simple components of a task (e.g., vertical and horizontal components of a cross; Shultz & Rivest 2001), and analogous tasks (e.g., vowel recognition of male speakers as relevant knowledge for recognizing female speakers; Rivest & Shultz 2002).

Because CC and KBCC networks start with little internal structure, they tend to learn simple ideas first and build more complex ideas on top of simpler ones, hence learning hierarchically. Associations between simple and complex representations are implemented in the connection weights between the network learning the complex task and a recruited sub-network having learned a simpler task. This addresses the second issue of P&V. Additionally, KBCC creates new connections when recruiting a sub-network into the current network topology. Only those new connections are trained, and therefore knowledge of the simpler task embedded in the recruited network is left intact.

There are two differences between KBCC and PARSER that should be further investigated. First, KBCC is a supervised learning algorithm, as opposed to PARSER, which is unsupervised. Using a variant of CC called recurrent cascade-correlation (RCC) (Fahlman 1991), it is possible to define the supervised task as the prediction of the next symbol in the input sequence, hence solving the problem of providing a target vector. Alternatively, if the task can be construed as an encoder problem, in which stimulus patterns are encoded onto hidden representations and then decoded onto output units, the need for a separate target vector also evaporates.

Second, KBCC does not learn all levels of a complex task as seamlessly as Perruchet and Vinter are aiming for. Instead, KBCC requires external intervention to present tasks in a timely and ordered fashion. After each sub-task is learned, KBCC is restarted with the next task as a goal and the networks trained on previous tasks as part of its candidate pool.

Nonetheless, KBCC does address the issues raised by P&V. Within KBCC, source networks are naturally formatted to serve as new coding primitives. This allows knowledge encoded in the connection weights of trained networks to be re-used to solve more complex problems. Furthermore, prior knowledge of simple tasks remains intact after recruitment, and new connection weights to the recruited network implement the association between simple and complex representations.

ACKNOWLEDGMENTS
Our research is supported by grants from the Natural Sciences and Engineering Research Council of Canada and from Quebec FCAR.

The self-organizing conundrum

Arnaud Destrebecqz and Axel Cleeremans
Cognitive Science Research Unit, Université Libre de Bruxelles, Fonds National de la Recherche Scientifique, B-1050, Belgium. **adestre@ulb.ac.be**
axcleer@ulb.ac.be **http://srsc.ulb.ac.be/axcWWW/axc.html**

Abstract: Perruchet and Vinter stop short of fully embracing the implications of their own SOC framework, and hence end up defending an implausible perspective on consciousness. We suggest instead that consciousness should be viewed as a graded dimension defined over quality of representation. This graded perspective eliminates the most problematic aspects of the cognitive unconscious without denying its existence altogether.