Self-organising navigational support in lifelong learning: how predecessors can lead the way.

José Janssen*, Colin Tattersall, Wim Waterink, Bert van den Berg, René van Es, Catherine Bolman, Rob Koper.

Open University of the Netherlands, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands

Abstract

Increased flexibility and modularisation in higher education complicates the process of learners finding their way through the offerings of higher education institutions. In lifelong learning, where learning opportunities are diverse and reach beyond institutional boundaries, it becomes even more complex to decide on a learning path. However, efficient and effective lifelong learning requires that learners can make well informed decisions. Drawing on principles of self-organisation and indirect social interaction, this article suggests solving the problem by analysing the paths followed by learners and feeding this information back as advice to learners facing navigational decisions. This article starts by introducing the principles of self-organisation and indirect social interaction. It describes how we expect the use of indirect social interaction using collaborative filtering to enhance effectiveness (completion rates and amount of progress) and efficiency (time taken to complete) in lifelong learning. The effects were tested in a controlled experiment, with the results showing effects on effectiveness though not on efficiency. The study shows that indirect feedback is a promising line of enquiry for navigational support in lifelong learning.

Keywords: distributed learning environments, lifelong learning, navigation

1. Introduction: the need for navigational support in education and lifelong learning

In general terms, navigation can be defined as “the process of determining a path to be travelled by any object through any environment” (e.g., Darken & Silbert, 1993). Several studies into student progress and retention highlight navigational issues in educational institutions. Yorke (1999) concludes that “As the unitization of curricula spreads through higher education, so there is a need for greater guidance for students to navigate their way through the schemes.” (p. 105). Research at the Open University of the Netherlands reveals that students feel a need for adequate information on further study possibilities in an early stage of their study and that they find it hard to gain an overview of the number of modules and the best sequence of study. Here information overload seems to cause the problem, rather than a lack of information (Joosten & Poelmans, 1998). Martinez and Munday (1998) point to “presentation of course/programme overviews” and “sequenced, structured course work of progressive difficulty” as part of the solution to the drop-out problem. Simpson (2004) mentions a more recent (2002) survey of students withdrawing from courses at the British Open University where 21% of the withdrawers identify “inadequate course choice guidance” as a reason for dissatisfaction.

Although findings from research indicate there is a relationship between navigation/planning problems and drop-out, they also reveal that it is only one factor among many others (Bean & Metzner, 1985; Kember, 1990; Rovai, 2003). Unfortunately, although research in the field identifies factors such as study-advising and program fit as of influence on retention, few clues as to the nature of advice are available (Martinez & Munday, 1998; Chyung, 2001; Rovai, 2003). Simpson (2004) suggests several alternatives to costly one-to-one advice: diagnostic materials, ‘taster’ materials and student views - all with their own limitations. This leads the author to conclude that they should probably be used in combination, although this may prove too burdensome for students.

* Corresponding author: José Janssen. E-mail: jose.janssen@ou.nl Fax: +31455762802

At present, faculties of the Open University of the Netherlands recommend a certain route through the courses available in their programmes. To some extent there is a ‘natural order’ (i.e. the contents of one course require prior knowledge offered in another), but apart from these kind of interdependencies, the recommendations hold little “empirical base”. That is to say, they are not guided by knowledge of actual sequences students (prefer to) follow and/or the extent to which they proved to be successful. Useful though these recommendations may be, and responsive to learner’s need for consistency and clarity of programs, they run counter to principles of learner centeredness and learner control. The need for alternative solutions to pre-planned routes is even more pressing in lifelong learning where learning opportunities are more diverse and reach beyond institutional boundaries. The concept of Learning Networks (Koper, Rusman & Sloep, 2005) provides a framework for addressing this complexity. Learning Networks (LNs) are self-organised, distributed eLearning systems designed to facilitate learner controlled lifelong learning in particular knowledge domains. Self-organised here means that organisational structures evolve from the actions and interactions of individuals, rather than being pre-defined; bottom up rather than top down. An important motive for bringing about self-organisation in Learning Networks lies in increased efficiency of the support structure (Koper, Giesbers, Van Rosmalen, Sloep, Van Bruggen & Tattersall et al., in press). Figure 1 is a simplified representation of a Learning Network in a certain domain D. The Network contains Activity Nodes (ANs, learning events or units of learning) which have to be mastered for the attainment of a certain objective or competency level. These activity nodes are the offerings of different educational providers.

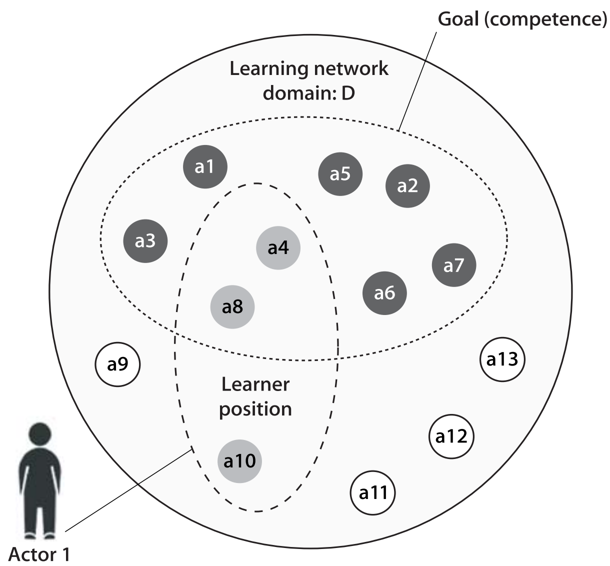

Fig. 1. Illustration of the concepts Learner position and Goal in a Learning Network.

The learner’s goal is the level of competence he or she would like to achieve. A route consists of one or more ANs that lead to the achievement of that level of competence. A “To do list” gives the ANs that still need to be completed. The learner in figure 1 has to complete six more ANs to reach the goal: the dark ANs in figure 1. The white ANs in the figure are out of scope for this learner at this stage: they have not been mastered yet (are not part of the learner’s position) nor constitute part of the goal. The grey ANs represent the learner’s position: either ANs that have been completed within the current Learning Network, or ANs accredited through prior learning. As the learner proceeds through the Learning Network, working towards his or her goal completing one AN after another, a learning track is built up, consisting of the sequence of ANs the learner has completed.

This is a simplified figure; in reality there may be many more ANs perhaps including alternatives from which to choose. Even this simplified example raises the question of how best to work toward the goal; what path should be followed through the To-do-list. Alternatives to one-to-one advice and pre-planned routes for navigational support can be sought in several directions (Tattersall, Manderveld, Van den Berg, Van Es, Janssen & Koper, 2005). Social navigation, like the student views proposed by Simpson (2004), is one of the alternatives. However, as Nichols (1997) points out, social filtering systems using explicit ratings require a large number of ratings to remain viable and users might consider it too much of a burden to rate ANs. A way to avoid placing a burden on learners is to rely on indirect social navigation, a concept closely related to the principle of self-organisation.

2. Self-organised indirect social navigation

An example often used to explain the concept of self-organisation is that of ant colonies, where individual ants leave behind pheromones (chemical traces) that signal to other ants the shortest routes to food (Bonabeau, Dorigo & Theraulaz, 1999). These trails feed information back on the progress of preceding ants. For self-organisation to occur in a network, actors have to have a high level of interactivity as well as access to feedback concerning the performance of similar others ('neighbours') in the network (Koper, 2005). This does not necessarily require *direct* interaction or feedback, but might take place through indirect feedback, also known as stigmergy: traces left and modifications made by individuals in their environment can function as feedback (Theraulaz & Bonabeau, 1999). Where Rovai (2003) states that "other students, staff, and faculty may not be readily accessible that can provide students with the information that they seek" our approach to offer indirect feedback might help bridge the gap: other students may be consulted as a source of information, albeit indirectly, by offering information on their navigational choices: the traces they have left behind while working towards their goal within the Learning Network.

Our study explored the use of this principle of stigmergy in offering wayfinding support, aimed at increasing the effectiveness (i.e. producing the desired effect) and efficiency (i.e. producing the desired effect with a minimum of effort) of Learning Networks. More specifically, we offered learners feedback concerning the choices and results of preceding learners aiming for the same goal. Our approach to offer navigational support in a Learning Network based on choices made by those who went before is quite similar to collaborative filtering used in recommender systems, where knowledge about the preferences of a large number of users is used to recommend items to a single, presently active user (Pennock & Horvitz, 1999). Our approach exploits information on choices/actions of users to derive (calculate) a recommendation. There are various types of information that could be offered as feedback to learners: information on the fastest route, the route with highest success or satisfaction rates, or a combination of several of these leaving it to the student to choose between these options. In order to feedback this information, a collective log of learner interactions within the Learning Network is filtered and processed as described in Tattersall et al. (2005). In our study, learners were offered feedback regarding the best next step, based on the number of times an AN had been *successfully completed.* In the study, an AN was successfully completed when a learner passed the assessment related to the AN. A similar approach is followed in work carried out for the French e-learning company Paraschool (Semet, Lutton & Collet, 2003; Valigiani, Jamont, Bourgeois Republique, Biojout, Lutton & Collet, 2005), although the feedback we propose is independent of any predefined or preferred routes. The feedback in this study is calculated as follows: if an activity node AN1 has been completed by 10 learners and 4 of those learners went on to successfully complete AN4, whereas 2 went on to successfully complete AN3, the advice for the next best step to a learner who has just completed AN1 as a first node, will be a random draw from the set {AN4,AN4,AN4,AN4,AN3,AN3}. Taking a random draw ensures that the most frequently completed AN is most likely to be recommended, while leaving room for other successfully completed ANs to be recommended as well, thereby avoiding sub optimal convergence to a single next step. For a more detailed explanation of this feedback calculation and the rationale behind it see: Koper (2005).

We expect that the navigation tool will enhance effectiveness and efficiency in Learning Networks since navigational support will facilitate planning decisions and reduce the risk of information overload by offering accessible and more learner centred (i.e. related to learner's present position) planning information. Moreover, as the feedback makes use of success rates, we expect learners to make better choices based on "tried and tested" sequences.

The nature of distance education and lifelong learning and, more generally, discussions on definitions and calculations of output and dropout in education (Cookson, 1990; Fritsch, H. 1991; Kember 1995; Reimann, 2004; Woodley, de Lange & Tanewski, 2001; Yorke, 1998) suggest that by simply defining effectiveness in terms of goal attainment we would be overlooking the fact that progress may have been made despite non-completion. In our study we will therefore not only look at goal attainment (the number of learners achieving a predefined goal), but also at the amount of progress made (the number of ANs that have been completed). Efficiency on the other hand will be indicated by a single variable: the time it takes to attain the goal.

The following hypotheses were tested in an experiment, using a feedback tool recommending a best next step based on successful choices of other learners:

1. *Offering feedback on the best next step, based on past choices of successful learners will result in increased effectiveness as indicated by both the amount of progress made (the number of ANs completed) and goal attainment (the proportion of learners reaching a predefined goal).*
2. *Offering feedback on the best next step, based on past choices of successful learners, will result in increased efficiency as indicated by the time required to attain the goal.*

The method section describes the experimental design in more detail as well as the way the advice regarding the best next step was presented to learners.

3. Method

To test the assumed effects of the navigational feedback a true experiment (Ross & Morrison, 1996) was carried out in which participants were randomly assigned to an experimental group that was offered feedback and a control group that proceeded through an otherwise identical Learning Network without any feedback.

3.1. Participants

Participation in the LN was free, i.e. no fees were charged and a popular topic was chosen for the Learning Network, namely the Internet. The target group was defined as people who have some experience with Internet - surfing the worldwide web and using email - and who face questions such as: How safe is it to buy things on the Web? Are there more efficient ways to search the Web? What do I need to do to ensure that my children are not confronted with ‘adult’ websites or adverts on the Web?

The recruitment announcement highlighted that the course[1] was designed with the purpose of testing new technology, that it would take approximately 22 hours to study the course, that the course would be available for three months starting in March and that completion of the course would be rewarded with a certificate. Prerequisite knowledge was defined as: “having some experience with Internet (surfing the web and using email) and a passive understanding of English”. At the start of the course participants were asked to fill in a small questionnaire aimed to gather some basic background information about the learners: age, sex, educational level and computer skills.

A group of 1011 people initially showed an interest in taking the course. They were randomly assigned to either experimental group or control group and given login details accordingly. Twenty percent (n=203) did not log into their assigned Learning Network site, and this group of non-starters is not included in our study of the effects of navigational support. This leaves a group of 808 learners who did enter the Learning Networks; 398 in the control group and 410 in the experimental group. Response rates on the background variables questionnaire were about 60%, showing that overall there were more women (59%), people over 45 years old (57%) and people with a high educational level (higher professional education or university level; 63%). Finally, 48% said their computer skills were poor or very poor.

3.2. Materials

A Learning Network was designed with the purpose of creating an appropriate experimental context to present and test the effects of navigational feedback. Designing the Learning Network, we took into account that:

- Due to time constraints the experiment should take no more than three months, meaning that learners must be able to reach the goal within three months.
- The Learning Network should contain sufficient ANs, so a navigational “problem” does indeed present itself.
- Completion of an AN must be “formally” established so that learning tracks can be determined and feedback can be derived from them.
- The Learning Network should reflect as closely as possible a realistic lifelong learning context, being both intrinsically and extrinsically motivating (reaching the goal would be rewarded by a certificate).

[1] In communications with learners we used the more familiar concepts “course” to indicate the entire Learning Network and “course part” for a single AN.

Eleven ANs were developed with the following titles: “The many roads to the internet”, “Web searching”, “Chatting“, “Secure payments on the internet”, “Do more with Internet Explorer”, “Worms and Horses”, “Beating spam & spy ware”, “Interesting and pleasant sites”, “Watching and listening on the internet”, “Dealing with inappropriate web content”, “Making a personal web page”. The ANs were designed to take an average of two hours to complete. Formal completion of an AN was established through the use of a short test consisting of five equally weighted questions. A score of 60% or more indicated successful completion.

Two Learning Networks consisting of the above ANs were created in the open source learning environment Moodle (Dougiamas, 2004), one for the experimental group, the other for the control group. In Moodle, each AN was modelled as a separate entity, thus ensuring that the learning environment kept adequate log records needed to provide the feedback and to test the hypothesis. The learning environment was modified such that all learners, both in the control and experimental group, received an overview of the ANs in the Learning Network, with a list of completed ANs on the left hand side and a To do list on the right hand side. For learners in the experimental group the overview also contained the advice: “Continue with: [the best next step, based on successful choices of other learners]”. Figure 2 shows the overview for a learner in the experimental group.

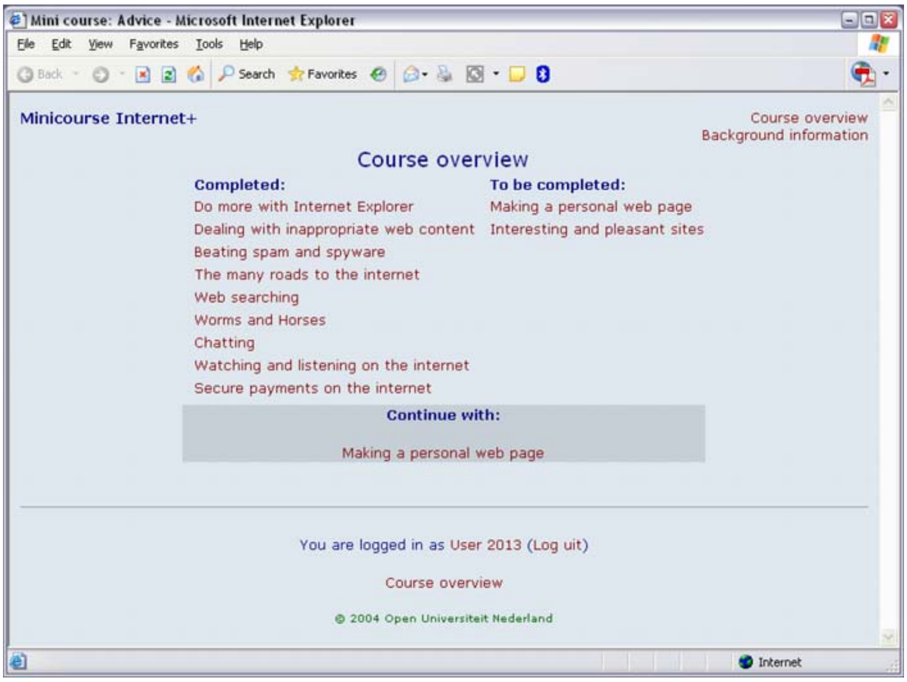

Fig. 2. Overview for a learner in the experimental group

The only difference in the experimental set-up for the control group lies in the absence of the “Continue with” area. The order of the ANs in the To do list was reshuffled each time the page was viewed so that there would be no effect in the sequencing of ANs due to the presentation in a fixed list.

3.3. Procedure

Participants were randomly assigned to an experimental group and a control group. Both groups received login details for their respective websites and a link to further instructions in an on-line user manual, and were informed that certificates would be issued to learners completing all eleven ANs in the experimental time period. Participants could study the ANs in any order, though learners in the experimental group were advised to follow the recommendation “Continue with:”.

All participants were told the list of ANs to complete would appear in a different order each time, but were not told why this was the case. It was explained that they would be randomly assigned to one of two groups who would work in a slightly different environment but with the same course content. There were separate email helpdesks for both groups offering technical and practical support. We deliberately chose not to offer any support regarding the course contents as this might have affected the experiment. During the three months the course was running, three newsletters were sent to inform students about technical topics that were raised via the helpdesk, and to remind them of the closing date of the course. The newsletters were identical for both groups and were sent simultaneously. The first newsletter was sent within a week as a number of learners had problems logging in to the course websites, and consequently turned to the email helpdesk for assistance. This first newsletter focused on those problems by explaining how to avoid mistakes with username and password and how to adapt cookie and internet security settings. A second newsletter was sent one month after the first and a final newsletter was sent as a last reminder of the closing date, ten days prior to the end of the experimental period.

3.4. Analyses

The first hypothesis stating that the feedback offered will result in more effective life long learning was tested for two different variables: amount of progress (the number of ANs completed over time) and goal attainment (the proportion of learners receiving a certificate; i.e. completing all 11 ANs). Goal attainment was measured by a single indicator, namely the proportion of learners having completed all 11 ANs at the end of the experimental period. Progress, in contrast, was measured over time, allowing for a comparison of the way progress developed in both groups using multivariate analysis of variance for repeated measures. In the experimental period of thirteen weeks, measures were taken at three weekly intervals, with the exception of the first measurement which was done after four weeks. The average number of ANs completed over these four successive moments in the experiment, was analysed by means of linear and quadratic trend analysis. Average ANs scores were transformed into linear and quadratic trend contrast scores by means of computation of orthogonal polynomials. Multivariate analysis of variance for repeated measures was applied on these contrast variables, which were chosen a priori, with Group (containing 2 different values: experimental or control) as between-subjects factor and Progress (4 values for four successive moments) as within-subjects factor. In case of significant interactions of contrast scores with Group or Progress, testing of simple contrast effects were performed. Due to the a priori character of these tests, they were performed with the conventional Type I error of .05 (cf. Tabachnick & Fidell, 2001).

In order to compare goal attainment for the experimental and control group a χ^2 test was used.

To test the second hypothesis concerning the effect on efficiency, a t-test was used to compare the average time taken to complete 11 ANs in both groups. The time taken to complete was measured by counting the number of days between initial login and completion of the final AN.

4. Results

4.1. Effectiveness

The results for effectiveness will be described separately for the two variables amount of progress made towards achieving the goal (the number of ANs completed) and goal attainment (the number of learners completing all 11 ANs).

4.1.1. Amount of progress

The amount of progress made by learners as indicated by the number of ANs completed in the course of time in both groups is represented in Figure 3. The overall completed ANs over time was denoted by a significant positive linear trend ($F(1,806) = 586.91$, $p < .001$) and a significant positive quadratic trend ($F(1,806) = 10.55$, $p < .001$). This means that the total group of participants has made significant curvilinear progress over time.

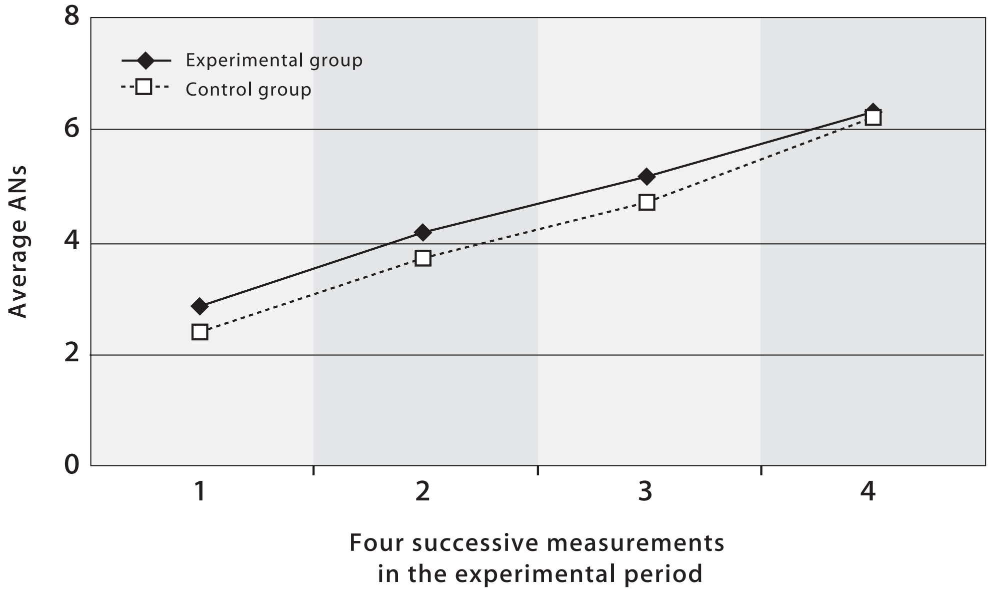

Fig. 3. Time course of progress (completed ANs) for the experimental and control group.

There was, however, no significant overall effect of Group, indicating that on average the two groups did not differ significantly. The interaction between Group and Progress was also not significant. However a significant effect of Group on the quadratic trend was found ($F(1,806) = 4.96$, $p < .05$), but not on the linear trend. Simple effects analysis showed only a significant linear increase for the experimental group (positive linear trend: $F(1,806) = 272.90$, $p < .001$) and a curvilinear increase for the control group (positive linear trend: $F(1,806) = 314.48$, $p < .001$; positive quadratic trend: $F(1,806) = 14.77$, $p < .001$). These results indicate that AN completion in the experimental group developed along a straight line, whereas in the control group the amount of progress made accelerated towards the end. Figure 3 indicates this: the average number of completed ANs is consistently higher in the experimental group except for the final measurement. In the end, the average number of completed ANs is about the same for both groups. This shift towards the end may have been influenced by an intervention, carried out ten days prior to the end of the experiment, when learners where reminded of the course deadline. To test the possibility that the intervention may have had an unintended and different impact for the control group, a repeated measurement analysis was performed for the last three weeks for learners who completed at least one AN.

Figure 4 shows study progress over the last three weeks of learners who completed one or more ANs during the experimental period.

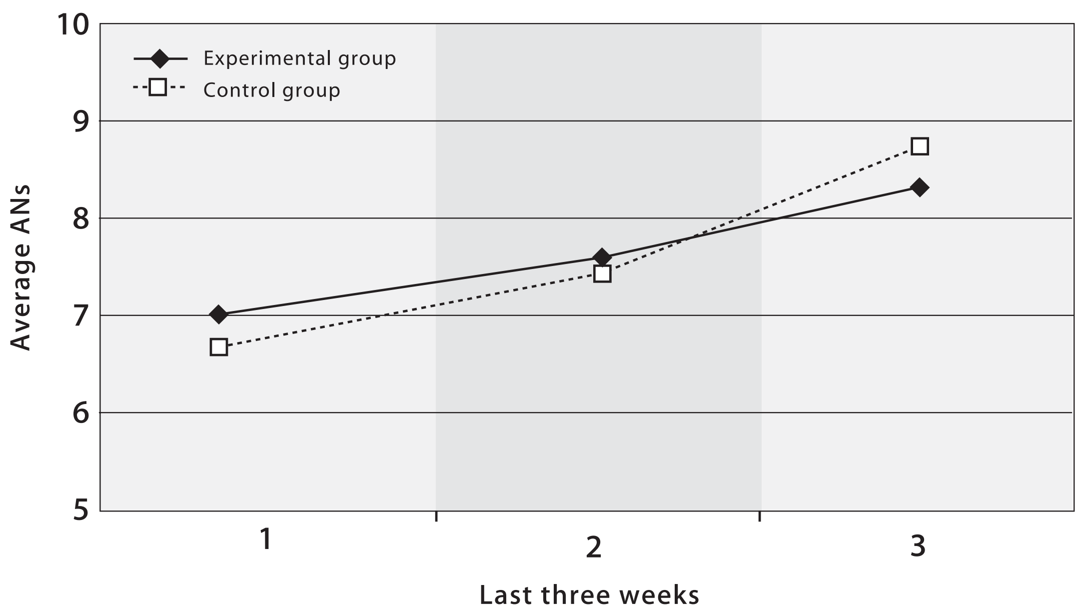

Fig. 4. Time course of the last three weeks of study progress of the experimental and control group.

The overall completed ANs over the last three weeks was denoted by a significant positive linear trend ($F(1,600) = 185.08$, $p < .001$) and a significant positive quadratic trend ($F(1,600) = 17.02$, $p < .001$). An overall significant effect of Group was not found. But there was a significant interaction between Group and Progress ($F(2,599) = 4.37$, $p < .05$) and there was a significant effect of Group on the linear trend ($F(1,600) = 8{,}67$, $p < .005$). Simple effects analysis showed a significant linear increase for the experimental group (positive linear trend: $F(1,600) = 59.79$, $p < .001$), and a curvilinear increase for the control group (positive linear trend: $F(1,600) = 130.43$, $p < .001$; positive quadratic trend: $F(1,417) = 15.80$, $p < .001$). This shows that the intervention indeed only had an effect for the control group. As a result, further analyses focused on the period up to the point where the intervention was made.

Repeating the analysis for four measurements during the period prior to the intervention shows a significant effect for Group ($F(1,806) = 4.32$, $p <.05$) on the number of ANs completed, indicating that the amount of progress made by learners in the experimental group was significantly higher over the period up to the intervention.

4.1.2. Goal attainment

Table 1 shows completion rates in the control group and experimental group immediately prior to the intervention. The percentage of learners completing all 11 ANs is significantly higher in the experimental group (40,2%) than in the control group (33,4%) ($\chi^2 = 4.04$, df = 2, $p < 0.05$).

Table 1
Completion rates (percentages) in control group and experimental group prior to intervention

Completion of 11 Ans	Group	
	Control[a]	Experimental[b]
No	66.6	59.8
Yes	33.4	40.2

[a] *n*=398
[b] *n*=410

4.2. Efficiency

For the group of learners who had completed 11 ANs at the point of intervention, the average number of days elapsed between enrolment for the first AN and completion of the 11^{th} AN was 36.49 in the experimental group, compared to 38.96 in the control group. Although learners in the experimental group reached the goal in fewer days, a t-test comparing these means shows that this difference is not significant.

5. Conclusions and discussion

The results of the experiment lead us to conclude that our approach to navigational support based on feeding back the choices of successful learners enhances effectiveness, though not efficiency, in lifelong learning. Improved effectiveness was not clear from the initial analysis. However, subsequent analyses corrected for the unexpected and unequal effect of the course deadline reminder and showed a significantly higher amount of progress and higher completion rates in the experimental group. The navigational support proposed in this study did not have a significantly positive effect on efficiency, i.e. the time taken to complete 11 ANs.

There are, however, a number of limitations with the experiment. First, although our work addresses lifelong learning, the limited experimental period of three months inevitably excludes some of the navigational and motivational problems faced by lifelong learners; a study of several years would be required to better reflect the intended application of our approach. A second limitation lies in our use of elapsed time rather than actual study time to indicate the time taken to complete 11 ANs. The use of this rather crude measure may mask significant differences in efficiency between the groups; subsequent work would benefit from a more accurate measurement of study time, although this is fraught with difficulties. Third, the experimental set-up did not force learners to take the recommended next step, and we do not know to what extent learners actually followed the advice. This resulted from the dynamic, just-in-time nature of the recommendation which was recalculated each time the overview page was refreshed. As a result the improved effectiveness can not be unambiguously ascribed to the recommendation itself; the mere presence of an advisory aid may have stimulated the experimental group. An additional experiment involving a control group receiving “fake” advice would help clarify this point. A further clarification of the results could be reached by investigating the extent to which the advice is followed and how this relates to the effects identified in this experiment. Finally, the best next AN was calculated using the most recently completed AN and relating this to the ANs successfully completed next by predecessors (but not yet completed by the learner). Extending the calculations to include a greater proportion of the position of a learner (rather than only the most recently completed AN) or even the full track of a learner and to move beyond ‘next best single step’ to the advising of next best sequences, might lead to stronger effects.

Further research is needed to address these limitations and to reveal whether alternative feedback calculations would have a greater impact on effectiveness and efficiency in lifelong learners. Alternatives to the feedback presently offered (based on frequency of success) include using study time, popularity or final grade. In addition, learner characteristics such as age, sex or competence level could be taken into account to filter the data before calculating the feedback, leading to recommendations which would allow the next best step taken by women, undergraduates, or the over fifties to be presented.

Despite the limitations of the present study, we believe it shows that the use of self-organisation principles offers a promising line of enquiry for efficient and effective navigational support in lifelong learning.

References

Bean, J. & Metzner, B. (1985). A conceptual model of nontraditional undergraduate student attrition. *Review of Educational Research, 55,* 485-650.

Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). Introduction. In E. Bonabeau (Ed.), *Swarm Intelligence*. Oxford: Oxford University Press.

Cookson, P.S. (1990). Persistence in Distance Education: A Review. In M.G. Moore (Ed.), *Issues in American Distance Education.* Pergamon Press.

Chyung, S.Y. (2001). Systematic and Systemic Approaches to Reducing Attrition Rates in Online Higher Education. American Journal of Distance Education, Vol. ?5, No.3.

Darken, R.P., & Silbert, J.L. (1993). A toolset for Navigation in Virtual Environments. *Proceedings of ACM User Interface Software & Technology*, 157-165.

Dieberger, A. (2003). Social Connotations of Space in the Design for Virtual Communities and Social Navigation. In K. H. D.Benyon (Ed.), *Designing Information Spaces : The Social Navigation Approach*. London: Springer.

Dougiamas, M. (2004). *Moodle*. Retrieved online 27/02/2004 at: http://moodle.org/

Fritsch, H. (1991). Drop In and Drop Out: The Need of Definition. In B. Holmberg & G.E. Ortner (Eds.), *Research into Distance Education.* Frankfurt am Main: Peter Lang.

Joosten, G., & Poelmans, P. (1998). Follow up van de eerste kennismakers: studievoortgang, -plannen en doorstroom. Heerlen: Open University of the Netherlands

Kember, D. (1990). The use of a model to derive interventions which might reduce drop-out from distance education courses. *Higher Education*, 20, 11-24.

Kember, D. (1995). *Open Learning Courses for Adults. A model of student progress.* Educational Technology Publications, New Jersey: Englewood Cliffs.

Koper, E.J.R., Giesbers, B., Van Rosmalen, P., Sloep, P., Van Bruggen, J., Tattersall, C., et al. (in press). A Design Model for Lifelong Learning Networks. *Interactive Learning Environments*.

Koper, E.J.R. (2005). Increasing Learner Retention in a Simulated Learning Network Using Indirect Social Interaction. *Journal of Artificial Societies and Social Simulation, 8* (2). Retrieved online 04/05/2005 at http://jasss.soc.surrey.ac.uk/8/2/5.html

Koper, E.J.R., Rusman, E., & Sloep, P. (2005). Effective Learning Networks. *Lifelong Learning in Europe, 1*, 18-27.

Martinez, P., & Munday, F. (1998). 9,000 Voices: student persistence and drop-out in further education. FEDA-report, Vol. 2, No 7. Retrieved online 02/12/2003 at http://www.lsda.org.uk/files/PDF/ISSN14607034-3.pdf

Nichols, D.M. (1997). Implicit Rating and Filtering. In *Proceedings of the 5th DELOS Workshop on Filtering and Collaborative Filtering*, 10-12 November 1997, Budapest, Hungary. Retrieved online 31/05/2005 at http://www.comp.lancs.ac.uk/computing/research/cseg/projects/ariadne/docs/delos5.html

Pennock, D. M., & Horvitz E. (1999). Collaborative Filtering by Personality Diagnosis: A Hybrid Memory- and Model-based Approach. In *IJCAI Workshop on Machine Learning for Information Filtering*, International Joint Conference on Artificial Intelligence, August 1999, Stockholm, Sweden. Retrieved online 31/05/2005 at http://www.research.microsoft.com/~horvitz/cfpd.htm

Reimann, N. (2004). Calculating Non-completion Rates for Modules on Institution-wide Language Programmes: some observations on the nature of seemingly objective figures. *Journal of Further and Higher Education, 28 (2),* 139 -152.

Richey, R., & Nelson, W. (1996). Developmental Research. In: Jonassen, David H. (ed.), *Handbook of Research for Educational Communications and Technology: A project of the Association for Educational Communications and Technology*. Macmillan Library Reference. 1213-1245

Ross, S. & Morrison, G. (1996). Experimental Research Methods. In: Jonassen, David H. (ed.), *Handbook of Research for Educational Communications and Technology: A project of the Association for Educational Communications and Technology*. Macmillan Library Reference. 1148-1170

Rovai, A. P. (2003). In search of higher persistence rates in distance education online programs. *Internet and Higher Education, 6*, 1 – 16.

Semet, Y., Lutton, E., & Collet, P. (2003). Ant Colony Optimisation for E-Learning: Observing the Emergence of Pedagogic Suggestions. Paper presented at IEEE SIS'03: IEEE Swarm Intelligence Symposium, April 24-26, 2003, Indianapolis, Indiana, USA.

Simpson, O. (2004). Access, Retention and Course Choice in Online, Open and Distance Learning. Paper presented at the Third Eden Research Workshop, Oldenburg, Germany. Retrieved online 05/07/2005 at http://www.eurodl.org/materials/contrib/2004/Ormond_Simpson.html

Tabachnick, B., & Fidell, L. (2001). *Using Multivariate Statistics* (4th edition). Boston: Allyn and Bacon.

Tattersall, C., Manderveld, J., Van den Berg, B., Van Es, R., Janssen, J., & Koper, R. (2005). Self Organising Wayfinding Support for Lifelong Learners. *Education and Information Technologies, 10* (1-2), 111-123.

Theraulaz G., & Bonabeau E. (1999). A Brief History of Stigmergy. *Artificial life, 5*, 97-116.

Valigiani, G., Jamont, Y., Bourgeois Republique, C., Biojout, R., Lutton, E., & Collet, P. (2005). Experimenting with a Real-Size Man-Hill to Optimize Pedagogical Paths. Paper presented at the 2005 ACM Symposium on Applied Computing, March 13-17, Santa Fe, New Mexico, USA. Retrieved online 10/08/05 at http://portal.acm.org/ft_gateway.cfm?id=1066683&type=pdf

Woodley, A., de Lange, P., & Tanewski, G. (2001). Student Progress in Distance Education: Kember's model re-visited. *Open Learning, 16* (2), 113-131.

Yorke, M. (1998). Non-completion of full-time and sandwich students in English higher education: costs to the public purse and some implications. *Higher Education, 36*, 181-194.

Yorke, M. (1999). *Leaving Early. Undergraduate Non-completion in Higher Education.* London: Falmer Press.