Vocabulary Learning and Instruction
Volume 11, Issue 1, December 2022
http://vli-journal.org

Vocabulary Networks Workshop 2: Activating Words in a Network

Paul Meara[a] and Imma Miralpeix[b]
[a]*Swansea University;* [b]*Universitat de Barcelona*

Abstract

Workshop 2 explores some larger vocabulary networks than those presented in Workshop 1 (1,000 words instead of 100), and introduces some additional ways of activating words in a network (e.g., by using new parameters in our models). The simulations will show that raising a network's activation level may not be as easy as we would have expected, and that different kinds of events may cause either temporary or permanent changes in the networks. Findings will be discussed in relation to the previous literature on passive and active vocabularies. Several implications for vocabulary learning and teaching will also be considered.

Keywords: Active vocabulary, Attractor states, Network structures, Simulations, Vocabulary modeling

1 Introduction

In the first of these workshops (Meara & Miralpeix, 2022), we looked at some very simple models of a vocabulary network. The models were small—only 100 words—and the words they contained had only a minimal set of properties. Each word was linked to two other words, and each word responded to input from these linking words in a fixed way. The simulations showed that a simple network of this sort behaves in a way that is more complex than we might have predicted, and we saw that some interesting emergent features are beginning to appear even in these tiny models. In particular, we noted that the networks naturally and quickly move into a stable attractor state. We also noted that whether a word can be described as active or passive appears to be a property of the network as a whole, rather than an intrinsic characteristic of the word itself. In this workshop, we will explore these ideas in more detail.

2 Some Bigger Models

The models that we use in this section are essentially the same as the models we used in Part 1 of the Workshop, but they are much bigger. Here, the models consist of 1,000 words, rather than 100 — 10 times bigger than the networks we looked at in Part 1. This is still small in comparison with a real vocabulary which might contain several thousand words, but the networks in Workshop 2 are not totally unrealistic: many second language (L2) learners have vocabularies about this size (cf. Nation, 1996 where the term "*Little Language*" is used to refer to the

Please cite this article as: Meara, P. M. and Miralpeix, I. (2022). Vocabulary networks workshop 2: Activating words in a network. *Vocabulary Learning and Instruction, 11*(1), 35–52. https://doi.org/10.7820/vli.v11.1.meara2

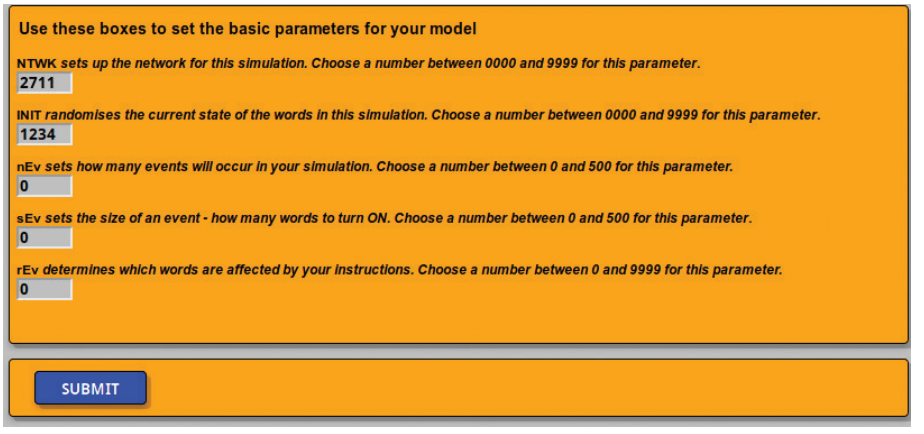


Figure 1. The Control Panel for Program-2.

first thousand words learners know, sufficient to make themselves understood and to perform a series of basic tasks). This means that we are beginning to look at vocabularies on a realistic scale.

In this workshop, we will continue using the programs we used in Workshop 1, so the basic approach will already be familiar to you. However, the bigger size of the vocabularies (1,000 words instead of 100) means that we need to change the way the programs display the information that they generate. Accordingly, we will start off by looking at the new display format, so that you can familiarize yourself with how it works.

To begin, you can go to the workshop home page: **https://www.lognostics.co.uk/Workshop/** and click on the button labelled **Program-2 Activating words**. The control panel for this program is shown in Figure 1.

This control panel is basically the same as the control panel in the previous Workshop, but here you have five variable parameters instead of two. The first two parameters, NTWK (network) and INIT (initialization), work in exactly the same way as they did in Part 1. NTWK controls the basic structure of the network: it determines the links between the words, and the way each word responds to its inputs. INIT determines the random activation pattern that the network starts out in. There are also three new parameters, all concerned with **events** that you can control. They allow you to make something happen to a network that has settled into a stable attractor state.

The next step is to run some simulations (you can run these simulations yourself) with the three new parameters set to 0. You do this by clicking the SUBMIT button. The program will return a report that looks something like Figure 2.

This new report screen shows you **the number of active words** in the network each time it updates itself, but it does not show the status of individual words—1,000 words is just too many for us to show them individually. We will be using lots of outputs like this throughout this Workshop, so it is important that you are able to interpret outputs like this one. In Figure 2, the network being examined is NTWK 2711. It is made up of 1,000 words. The number of active words in the

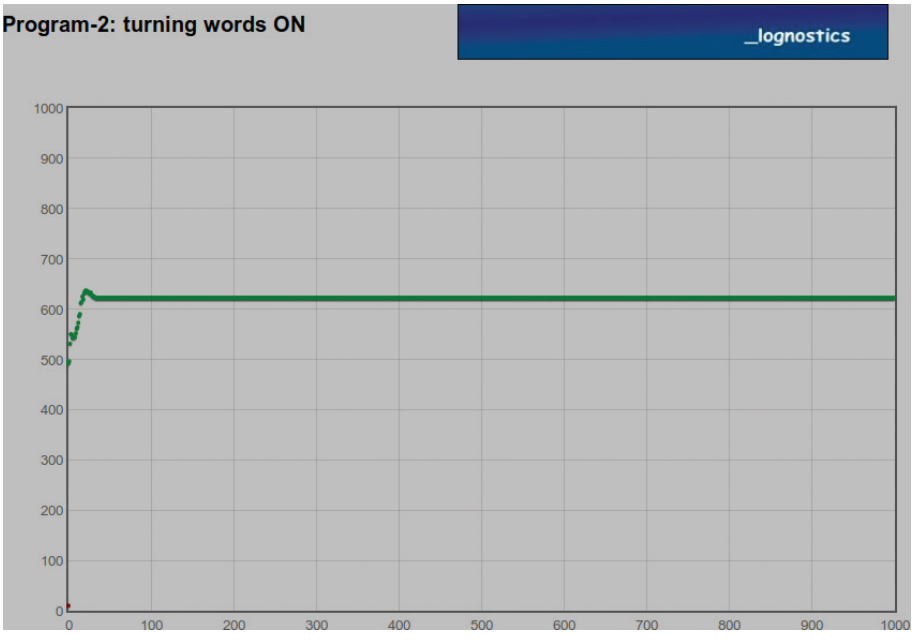


Figure 2. The Report Screen for Program-2 when NTWK: 2711, INIT: 1234, and the other Parameters are Set to 0.

network is indicated on the Y-axis by a line of green dots. The X-axis shows how this value varies over time as the network updates itself. Figure 2 shows that you have 1,000 updates in the simulation. In this report, the number of active words starts out at about 50% (recall that we initially set this value randomly, with about half of the words ON and half of the words OFF). Then the network adjusts its level of activation in response to this initial randomization. As the network updates itself, the number of activated words rises for a while, then falls back again, until after roughly 50 updates the network reaches an attractor state where about 620 words are activated. Once it reaches this attractor state, the network stays there. Experiment with different values for the NTWK parameter and the INIT parameter, and make sure that you understand what the new format is telling you.

Now let us build a simple real-world scenario around the data in Figure 2. Suppose we are dealing with a beginning learner of English—call her Maria. Maria, as a beginner learner, has a passive vocabulary of 1,000 words, but her active vocabulary is limited—only 620 words. The other words in her vocabulary—380 of them—are words that she knows only passively: these are words that she would recognize, but she is not able to produce them without help. Most research agrees that passive vocabularies tend to be bigger than active vocabularies, although there is no agreement on how large the possible gap between the two could be (e.g., see Melka, 1997 or Miralpeix, 2020 for a review), so these figures are arbitrary but plausible as a first stab. Obviously, this is not where we want Maria to be: we would really like to increase her productive vocabulary. Ideally, perhaps, we might like all 1,000 words to be part of Maria's productive vocabulary. The question we need to ask, then, is what do we need to do to the network in order

to achieve this outcome? The obvious answer is that we might be able to increase the number of active words in Maria's vocabulary just by turning more words ON, and hoping that doing so will result in a permanent increase in her active vocabulary size. Program-2 is designed to simulate this process.

You can go back now to the start page for Program-2. In addition to the two standard parameters, this start page includes three new parameters, **nEv, sEv,** and **rEv.** These three parameters allow you to control **events** that take place when you run your simulations. In Program-2, an event turns a number of randomly chosen words ON, and shows you what effect this has on the overall level of activity in the network.

- **nEv** sets the number of events you want to take place. For the moment, set the value of nEv to 1. This means that when the simulation runs you will get one event. However, you cannot control exactly when these events will take place, as the program decides this randomly.
- **sEv** sets the size of any events. This parameter controls how many words will be turned ON whenever an event occurs. Set the value of this parameter to 10. This means that each time an event takes place, 10 additional words will have their current value set to ON.
- **rEv** controls which words are affected by an event. You cannot control the fine detail of this in Program-2. For the moment, it is sufficient to know that specific values of rEv will always affect the same words each time you run the program. Set the value of rEv to 1234.

You can run the simulation by clicking on the **submit** button. You should get an output that looks something like Figure 3.

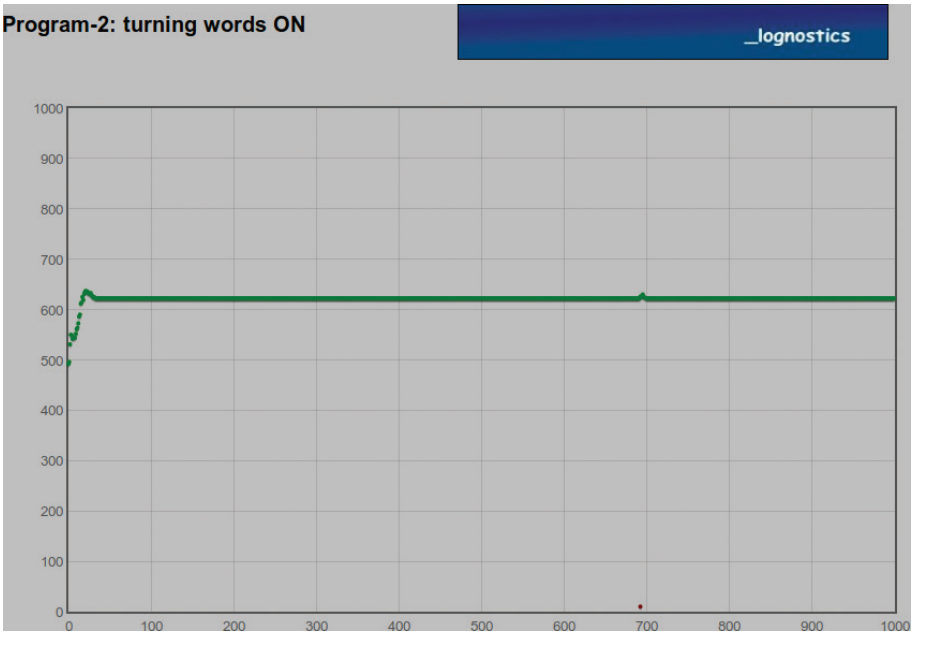


Figure 3. Program-2 Report page. NTWK: 2711, INT: 1234, nEv: 1, sEv: 10, REv: 1234.

Figure 3 tells us that this network starts off with about 500 words in their activated state and, after a small number of updates, the network settles into a stable attractor state. The red dot at the bottom of the figure tells us that a single event takes place around update 690. The effect of this event is to momentarily increase the number of activated words in Maria's vocabulary by 10 words. However, this burst of activation does not last very long—a small number of further updates brings the overall level of activation in the network back to its attractor state. So far, it looks as though just turning 10 words ON once does not have any long-term effect on Maria's vocabulary.

3 Program-2: Raising the Activation Level of a Stable Network

From here, you should be able to see that we can experiment with different combinations of the parameters, and examine the effects of turning words ON in different ways. You will perhaps be thinking of a number of obvious lines of enquiry and we have listed some of these below. The questions we suggest in this section do not have any "right answers." They are primarily designed to make you think about what the simulations are doing, and to suggest further questions that you might ask. What you should notice is that experimenting with different models makes you think about vocabulary networks in new ways. Here are some starter questions:

Are larger events more likely to have a permanent effect on the number of active words in a vocabulary network? You can test this idea by changing the value of the sEv parameter. For example, if you set sEv to 50, then the program will turn ON 50 words each time an event occurs instead of 10. How does NTWK 2711 react to a single event of this size? Increase the value of sEv and watch how this change makes a difference to the way this network reacts.

Do all networks react in the same way as NTWK 2711? You can examine this idea by varying the value of the NTWK parameter. Choose a set of values that give you networks with different attractor states. How often do you find a value of sEv that results in a permanent change in the overall activation level of a network? Can you identify some networks that do not react like NTWK 2711? Why do you think this might happen?

Do frequent small events have a permanent effect on the number of active words in a vocabulary network? You can test this idea by changing the value of the nEv parameter. For example, if you set nEv to 50 then the program will give you 50 events instead of one. Most often, you will find that simulations with repeated events look something like Figure 4. This simulation shows the same network that we saw in Figure 3, but with an increased value for the nEv parameter. In Figure 4, nEv has been set to 25, giving us 25 events, each activating 10 words. Again, events of this size and frequency have a small effect on the overall activation of words in the network. Each event causes a small ripple of activation in the network, but these changes are not permanent, and when the events stop around update 900, the network quickly returns to its attractor state.

Program-2 does not give you any direct control over when the events occur. Usually, the events will be spaced out more or less evenly across the X-axis of the

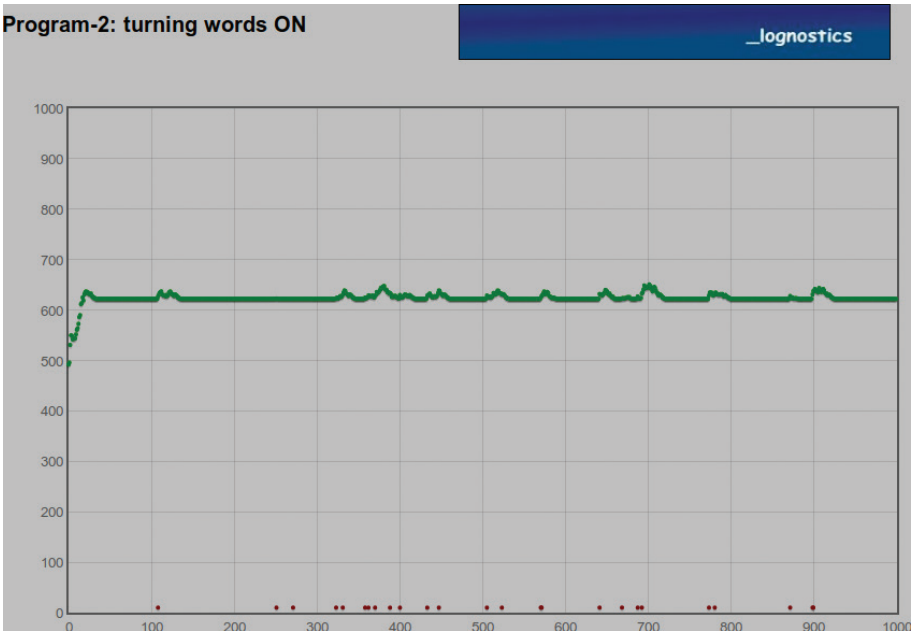


Figure 4. NTWK: 2711, INIT: 1234, nEv: 25, sEv: 10, rEv: 1234.

report page. However, if you experiment with different values of the rEv parameter, you will come across patterns of events that look more interesting. Sometimes you will get a cluster of events that happen in quick succession, and you might find that this causes more activation in the vocabulary network. This should lead you to ask:

Do clusters of small events have a permanent effect on the number of active words in a vocabulary network? You can test this idea by choosing a value of rEv that generates clusters of events. Generally speaking, you should find that small numbers of events which activate a small number of words have only a limited effect on the overall activation level of the vocabulary network, but when you increase the number of events and increase the size of each event, the overall activation levels will increase and be longer lasting. Sometimes you will find that the network will jump to a new, higher level attractor state, but occurrences of this sort are relatively rare.

You should find that there is a considerable amount of variation in the way the networks respond to different types of events. Figure 5 shows an example of how one network with a low level of activation behaves in response to a series of large input events. NTWK 1946 has a **very** small number of active words when it is in its attractor state, so we might expect it to be particularly responsive to events that activate large numbers of words. Think about how you would expect this network to behave before you go on to run the simulation yourself.

Here, we have asked for 20 events, and in each event 100 words are turned ON. The result is a huge amount of activity in the network, particularly when we have a cluster of events occurring in quick succession, such as the cluster around

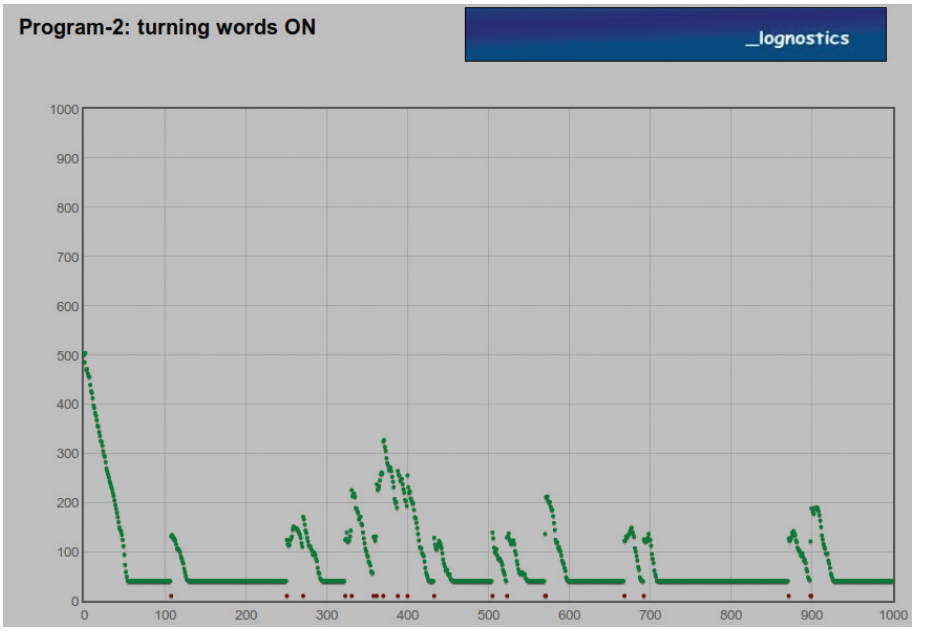


Figure 5. How a Network with a Low Level of Activation Responds to a Series of Activation Events. NTWK: 1946, INIT: 1234, nEv: 20, sEv: 100, rEv: 1234.

update 400. But surprisingly, none of these events seems to result in a permanent increase in the number of active words. This particular network stubbornly resists being shifted to a higher level of activation. Is this what you predicted? Why not? In fact, even if we ask for huge events, where 500 words are activated each time an event occurs, the long-term effect for this network appears to be negligible (cf. Figure 6).

However, not all models are as stubbornly resistant to change as this one is. Figure 7 shows an example of a model where a small number of input events does have a long-term effect on the overall activity level of the network. This figure shows a network that has an attractor state which already has a large number of active words. Each time an event occurs, the network moves to a slightly higher attractor state, and it seems to stay there.

This might prompt you to ask: **are highly active vocabulary networks more likely to respond to activation events than networks with low activity levels?** You can investigate this question by looking at a range of models with different attractor states. We recommend that you find a set of 10 networks with an attractor state with a lot of active words, 10 networks with an attractor state where about half the words are active, and 10 networks with an attractor state where only a few words are active. Does the level of activation in a network affect the way it responds to further activation events?

At this point in particular, you should be asking whether Program-2 is really a good model of how a vocabulary might become more active. The answer to this is that it probably is not a good model, but it is useful to think about why not. You

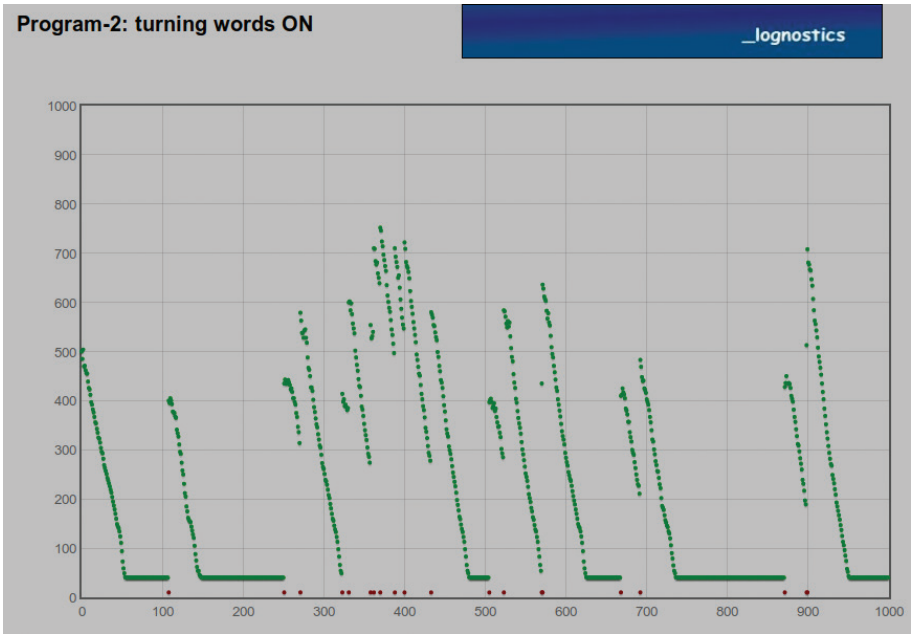


Figure 6. The Effect of Very Large Events (Each Event Turns ON 500 words). NTWK: 1946, INIT: 1234, nEv:20, sEv: 500, rEv: 1234.

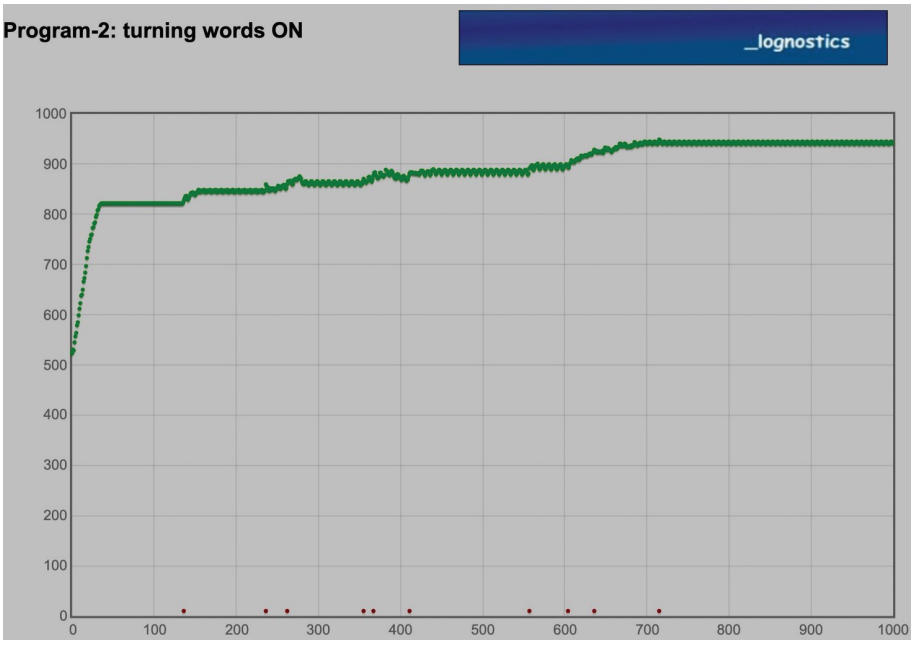


Figure 7. How a High Activation Network Responds to Large Input Events. NTWK: 1234, INIT: 1234, nEv: 10, sEv: 50, rEV: 500.

should also be thinking about what additional features would we need to build into the models to make them more realistic. Despite their lack of realism, are the models in this section perhaps drawing our attention to features of real vocabularies that our current theories do not take account of?

You should also be asking whether the models in this section are giving us the information that we really need. For example, is the total number of activated words in a vocabulary network the best way of assessing its performance? Maybe a different measure would be more useful? What might that be? When an event activates a large number of words, the overall activation level of the network rises instantly, but it fades away more slowly. This implies that there is a lot of underlying activity in the network that the report is not picking up on. Maybe, instead of just counting the number of active words at each update, we could look at how many words are changing their current state at each update of the network? Would this perhaps give us more insights into how the network is responding to the activation events?

More seriously, we might also ask whether it really makes sense to try to shift a network directly from one attractor state to another, just by activating a few words. All the simulations in this section have started off in one of their attractor states, and your explorations with Program-2 should have convinced you that it is surprisingly difficult to shift a network out of an attractor state into a new one just by turning words ON. However, it is possible that Program-2 is approaching this problem in the wrong way. Maybe vocabulary networks are not normally in an attractor state when events happen? We will look at this possibility with Program-3.

4 Program-3: Networks with a High Level of Temporary Activation

Program-3 lets you examine some of the more speculative ideas in this discussion. So far, we have only considered how we might shift a vocabulary network from one stable attractor state to another stable attractor state where more of its words are active. However, it is possible that networks in an attractor state are particularly resistant to change, and it is possible that a network that has been shifted out of its attractor state into an unstable state **(an excited state)** might be easier to manipulate. Basically, we need to explore whether an excited network can be prevented from falling back into its attractor state by activation events. Specifically, Program-3 lets you ask: **what happens in a vocabulary network that has been shifted out of its attractor state BEFORE any events to take place are allowed to take place**. Normally a network that has been shifted out of its attractor state would return to its attractor state after a small number of updates. However, it is plausible to imagine that a network that is in an excited state might respond differently. Perhaps randomly turning words ON in an excited network could prevent this return to normality, and bring the network to a different, higher attractor state? How effective would we expect this to be?

You can go now to the Workshop homepage and click on the Program-3 button. The start page for this program is shown in Figure 8: this page is similar to the start page for Program-2, but it contains an additional parameter that you can experiment with: **SPK (Spike).** The simulations in Program-3 run in exactly

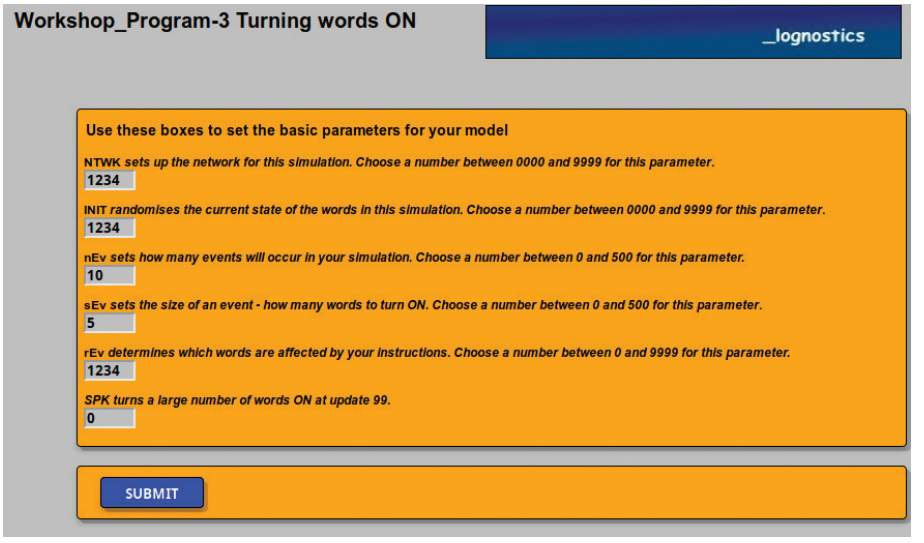


Figure 8. The Start Page for Program-3.

the same way as the earlier simulations in this section, but parameter SPK allows you to apply a single large spike of activity to the network, as well as a series of smaller input events.

This spike always occurs at update 99—the network will normally have settled into its stable attractor state by this time, so if you set the SPK parameter to a high value, it will jolt the network out of its steady state into a state which is more unstable. Maybe a small amount of input will be enough to keep the network in this state of heightened activity?

You can now set the NTWK parameter and the INIT parameter to 1003, nEv to 250, sEv to 200, rEv to 1003, and SPK to 500. This combination of parameters will give you a network which has a low-level attractor state (190 words are ON). The way this network behaves is illustrated in Figure 9.

This figure shows NTWK 1003 moving to its normal attractor state—just short of 200 active words (updates 0-99). At update 99 the program turns 500 words ON, and this results in a sudden uplift in the network's activation level. After just a few further updates, this figure has risen to just over 900 words: nearly every word in the network is now ON. Then, following this initial spike of activation, a series of events takes place, where each event turns ON 200 words. There are 250 events of this sort, but even when the events are as frequent and as large as this, the overall activation level of the network does not appear to be permanently affected. As long as the activation events continue, the number of active words in the vocabulary remains high. However, when the activation events stop at update 900, the network returns almost immediately to a low level of activity. Interestingly, this new attractor state is not exactly the same as the one this network started out with. The new attractor state has a few additional words that are permanently activated as a result of the massive amount of stimulation that we

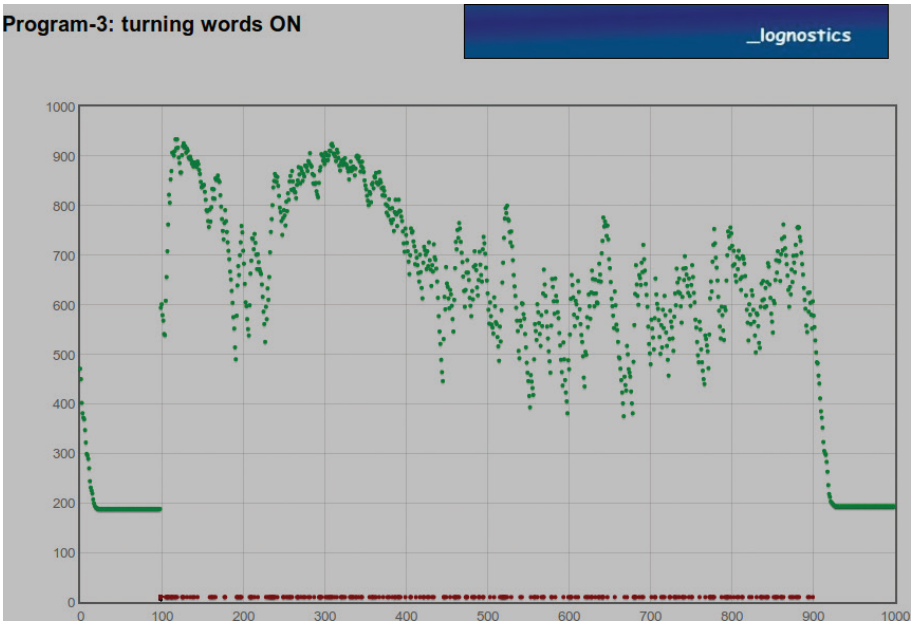


Figure 9. NTWK: 1003, INIT: 1003, nEv: 250, sEv: 200, rEv: 1003, SPK: 500.

have provided it with. This result suggests that turning words ON does sometimes cause a permanent change to the activity level of a network, but it is not obvious that a change of this kind is directly related to the amount of input that the network receives.

Here again, you should be asking questions about what Figure 9 is telling you. Would we get a bigger permanent uplift if we increased the number and the size of the events in this simulation? (Change the values of nEv and sEv to explore this question). Would we still get the uplift recorded in Figure 9 if we reduced the number of events, or if we reduced the size of the events? Does Program-3 produce simulations which exhibit threshold effects where a small change in the number of events or the size of an event generates a sudden change in the way the network behaves? Is the initial spike at update 99 really important, or is it just the combination of very frequent and very large events that gives us the pattern of activity reported in Figure 9? Do all networks with low levels of activity behave like this one? Do networks with a naturally high attractor state react in the same way as NTWK 1003? You can investigate all these questions by using Program-3 with different values for the parameters.

5 Discussion

Four main ideas seem to be emerging from the simulations in this section. First of all, the main feature that you can expect your simulations to show is that it is surprisingly difficult to shift a vocabulary network from one attractor level to different higher-level attractor. As we have seen with Program 2, a network in

an attractor state will usually just absorb small changes and go back to its stable resting state. The main exception to this is when your simulation implements a lot of events, and these events are large enough that they overwhelm the network's natural tendency to go back to its resting state. In this case, you can sometimes find that very high levels of activity can be supported by the network. However, cases where this happens appear to be rare. In most cases, a raised level of activity is only temporary. When the activation events stop, then the network will most often fall back to its stable resting state, even when the number of activated words has been very high.

What are the implications of this? It is difficult to draw any immediate conclusions from these simulations, because the parallels between the simulations and real life are a bit tenuous. We cannot turn passive words ON in real life—all we can do is provide some conditions in which some words might get turned ON, by getting students to read, for example, or sending them off on an intensive residential course. The simulations do not mirror treatments of this kind directly. Nevertheless, the inescapable conclusion seems to be that activities for learners that just turn passive vocabulary items ON temporarily are not in themselves going to be sufficient to increase students' active vocabulary on a permanent basis, and this **does** have some real-world implications. As Nation (1990, p. 94) noted, "if learners have a reasonably large receptive vocabulary but are unable to put enough of this into productive use, then the teacher needs to concentrate on activities that enrich the learning of known words and improve access to them."

Surprisingly, quite a lot of standard vocabulary exercises could be seen as events which turn passive words ON (see Meara, 1990). Consider, for example, an exercise like the one shown in Figure 10.

Suppose that our learner—Maria—does this exercise, as part of an activity designed to improve her vocabulary. If COMB is already part of Maria's active vocabulary, then this exercise just allows her to use a word she already knows: in the simulation, this would be equivalent to turning ON a word that is ON already, and we would not expect such an event to make any difference. If COMB is part of Maria's passive vocabulary, then she will recognize the word when she sees it, it will become temporarily active, and she will be able to use it correctly to fill the gap in the example. In the simulation, this would be the equivalent of turning ON a word that is currently OFF. But once the network is updated, it is unlikely that COMB will remain active: it is much more likely to revert to its normal passive state. The simulations in this chapter suggest, therefore, fill-in-the-gaps exercises like this may not be an effective way of making COMB an active vocabulary item for Maria.

Which word best fits the gap?
He is using a ____ to arrange his hair

1: razor 2: comb 3: tissue 4: toothbrush

Figure 10. A Typical Vocabulary Exercise.

This does not necessarily mean that exercises of the type illustrated in Figure 10 are no good. If you have used these exercises with your students in the past and found them effective, then a simple simulation like this one is not going to convince you that the exercise is not theoretically motivated and should therefore be abandoned. What should happen is that you realize this exercise is not working the way you think it is. The simulations seem to suggest that just activating a passive vocabulary item for an exercise is not enough to turn it from a Passive word into a permanently Active word, so we need to ask what **does** an exercise of this type do? Maybe the exercise works because it somehow makes COMB easier to activate, or maybe the juxtaposition of RAZOR, COMB, TISSUE, and TOOTHBRUSH changes the connections that COMB has with other words in the vocabulary network. These are ideas that we will explore in the next section of this workshop. What the simulations seem to confirm, though, is that this type of practice will not suffice to improve productive vocabulary use in tasks. In input-driven narratives, for example, where the speaker is pressured to be precise and find the exact wording (Skehan, 2009), words activated with exercises similar to that in Figure 10 may not be available, as they may not have become permanently active.

A second feature of these simulations is that large events which occur frequently are more likely to deliver a permanent change than small events that take place infrequently. This is not really surprising, perhaps—research has shown that immersion settings can speed up learning when compared to instructional settings, with less intense exposure (Foster, 2009; Milton & Meara, 1995; Zaytseva et al., 2018). Psychological studies have also indicated that an effective learning process is characterized by increasingly spaced repetitions, with a short gap in between early meetings and larger gaps in later meetings (Pimsleur, 1967). The simulations confirm these previous findings. Additionally, studies on reading or TV viewing have proved that retention rates are low if vocabulary is not recycled (Gesa, 2019; Waring & Tataki, 2003).

Note though, that the X-axis of the simulations does not actually specify how much time elapses between each update, so we cannot immediately convert the simulation into real time.

However, we could start to ask some questions which examine this idea. For instance, we might ask: in real life, for how long does a normally passive word remain activated, and for how long does the effect of this activation last? Clearly, if the effects of activation are all dissipated within a few hours, or even a few minutes, then it is going to be very difficult to meet the conditions necessary for a permanent uplift in the number of words in a learner's active vocabulary. In the simulation sets in this part of the Workshop, we have implemented 1,000 network updates which take place in a sequence, but we have not actually said anything about how often these updates take place. When does a network update take place: every few minutes, every day…? If the updates take place daily, then the 1,000 update events recorded in Figure 3 would take about 3 years to be completed, and this is probably not a realistic time scale. Maybe, then, we should be thinking about network updates that take place more than once a day—perhaps every hour. But this implies that temporary activation of a normally passive word would be very transient indeed—a word might be known now, but forgotten again in an hour's time. What are the implications of this?

The simulations in this section also assume that the network is updated regularly, but this might not be the case. We could have coded the program so that the network is updated at irregular intervals. Or we could have programed the simulations so that activation events do not always occur with an update event. This should make you think about whether updating the current status of ALL the words in the network simultaneously is the best way to model how a network responds to being stimulated. Maybe the network only partially updates itself? Maybe it only updates itself when it needs to? How would a network know when this was necessary?

Again, there are no obvious answers to these questions. The point of this discussion is to show you how running simulations makes you think about the assumptions that underpin normal practice, and to ask questions about these assumptions. The simulations do not provide immediate answers for these questions, but they do help to identify areas where current theory is underdeveloped, or in need of rethinking. They also provide a rich source of ideas for further research. Some straightforward projects which emerge from this simulation set would include the following questions, which do not typically appear in the current research literature:

In our simulations the vocabulary items can either be ACTIVE or PASSIVE, but a lot of the research literature assumes that this feature is graded rather than binary. Is this important? Would a network where words can have several different levels of activation, or different types of activation, behave very differently from the networks we have looked at here?

If you activate a passive vocabulary item, for how long does it remain active?

Do words differ very much in this respect?

Do different ways of activating passive vocabulary items produce longer or shorter activation periods for individual words?

Does vocabulary size make a difference? That is, does activation last longer if you have a larger vocabulary?

Is activation different for beginning language learners than for more advanced learners?

Is there a reliable method for identifying passive vocabulary items?

Are vocabulary networks normally in their stable activity state, or are they normally in an excited state where activity is greater than we would expect?

Do words in a network spontaneously (and randomly?) activate themselves? Does spontaneous activation of this sort have a noticeable effect on a network? How could we investigate this phenomenon in real language learners?

The third feature that should be emerging from your simulations is that the details of the model are more important than you might have expected. We already know that the network models in this section are only very rough approximations to real vocabularies, and it is important to stress that we are not suggesting that network models of this sort are intended to be accurate models of how real vocabularies work. Their main role is to be suggestive and provocative. But the same argument applies to the way we have modelled the events in this simulation set. Working

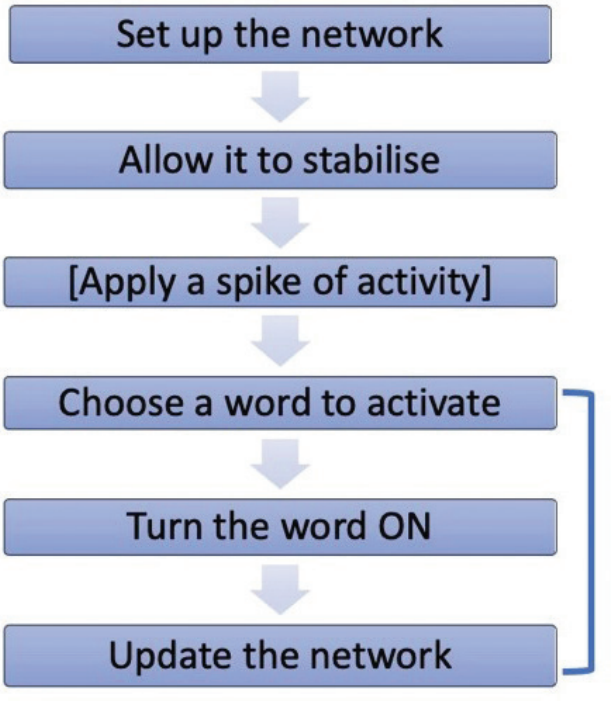


Figure 11. The Structure of the Models in Program-2 and Program-3.

with simulations forces you to think about whether the way you have modelled things is really sensible.

The basic approach we have adopted in this section is sketched out in Figure 11. Here, we have chosen to set up the network as a set of words which start off with fixed properties that do not change as the program runs. The only characteristic that changes in Program-2 and Program-3 is that some words move from their ON state to their OFF state, and vice versa. Their other properties remain fixed. The simulations do not change the way words are linked to other words, and they do not change the way words respond to other words. But maybe this approach is too rigid, and we should consider what happens in a network where these properties are not fixed? We will look at these issues in Part 3 of the Workshop.

We have also chosen to model networks that are of a fixed size—1,000 words in this section—but this means that we are not looking at dynamic networks which are growing or shrinking in size. We are also assuming that all words have the same characteristics and the same status. Our guess is that neither of these assumptions is correct, and a much safer assumption would be that words are much less homogeneous than we have modelled them. In particular, it seems a good guess that words which have only recently been learned do not perform in the same way as more well-established words. How could we build this factor into our models?

One major issue with the simulations in this Section is that we have made some assumptions that might have biased the way our vocabulary networks

behave. For example, we have chosen to activate words at random in the models in this section, but this has a number of important side-effects. One is that random activation will sometimes choose a word for activation that is already activated, and in that case the activation event will have no effect on the network. This is more likely to happen in a network where most of the words are already activated. For example, if 50% of the words in a network are already activated, then a randomly selected word has a 50% chance of being ON already, so half of our events will not make any substantial change to the activity level of the network. In a network where 90% of the words are already activated, only 1 event in 10 will produce an event that changes the network's activity level. This suggests that networks with low levels of activity ought to be more likely to show change than networks where there is a higher level of activation. The effectiveness of an event should get smaller as the activation level in the network increases. Does this matter? It is hard to tell, but our model clearly has a built-in bias in this respect, and we need to be aware of that. Of course, we could re-code the model so that it selected only those words for activation that are currently inactive. Would this make a difference to the way the models work? It is hard to tell without actually writing the code and testing it out. Again, randomly selecting a word for activation means that all words have an equal chance of being activated by an event, but we could rewrite the code so that some words were more likely to be activated than others. For example, we could program the simulations so that the first 50 words in each network are more likely to be activated than the remaining words. Would this make a difference to the way the models work? Probably. A number of researchers, notably Carter (1987) and Brezina and Gablasova (2015), have suggested that vocabularies typically include a set of core words which are more resistant to attrition than other words are. We will come back to this assumption in a later section of the Workshop.

Another assumption we have made in this section is that we have chosen to work with models which have been allowed to settle into one of their stable attractor states before any events occur. However, it is far from obvious that a vocabulary network will normally be in its stable attractor state. Maybe we should assume that, by default, a vocabulary network is NOT one of its attractor states. Rather, the norm might be for vocabulary networks to be in unstable, excited state, where lots of words are flipping between the ON and the OFF state? Maybe this is what makes them work? Maybe the whole idea that an ideal vocabulary is one where most of the words are permanently ON is wrong? Maybe a better model would involve building a vocabulary network which had a naturally low-level of activity but responded very strongly to any input by activating large swathes of words very quickly? Figure 9 shows a network that seems to be able to switch quickly between a state of low activity when it experiences a lot of external stimulation, and maybe this is an important feature of real vocabularies?

Finally, we need to ask whether simply activating a word is a good way of modeling what happens when teachers provide exercises aimed at consolidating vocabulary. In the models in this section, a word that has been activated is treated just like any other word in that it will revert back to its "natural" state immediately if the activation state of other words does not allow it to remain activated. But we could argue that newly activated words are more resilient than this. Let us say that a newly activated word stays active for a few updates before it reverts to

its normal activity level. That would mean that an event that activates five words would effectively activate them for the next few updates as well. The effect of this would be enhanced when two or three events occur in close proximity to each other, and that would mean that the timing of events becomes a critical feature of the way these models work.

The important thing here is that working with the detail of the models has thrown up a number of considerations that were not apparent when we began, and we have ended up with a lot more questions than we started out with. This is typically what happens when you work with formal simulations. Every decision that you make in the programming has knock on effects on the way the simulations behave.

The obvious conclusion is that you need to be aware of the shortcomings of the simulation programs. Furthermore, you need to be prepared to develop new simulations that explore the implications of these shortcomings. It is worth pointing out that this is a very different way of doing research. Traditional "real-world" research is mainly difficult because it is hard to assemble large groups of participants and test them effectively. In simulation research, you do not need hundreds of participants: you just run the simulations many times with different parameter values. In real-world research, if your work throws up something that looks intriguing or suggestive, it is very unlikely that you will get the chance to repeat your study and collect the data you need to follow it up. With simulation research, repeat studies are very easy to do. This is important, because it often allows you to eliminate some unproductive lines of thought without going to the trouble of running actual experimental studies. The simulations in this section, for example, strongly suggest that a study that teaches students a small number of words in a one-off treatment is unlikely to produce lasting effects on their productive vocabulary knowledge. Looking at how often studies of this kind crop up in the current research literature should give you pause for thought.

6 Conclusion

We hope that this preliminary excursion into simulations and modeling will have shown you that this approach has a lot to offer to vocabulary researchers. Even a very simple model raises lots of questions, which would not often arise in the context of a more traditional research approach. In this section we have focused on the idea that increasing an L2 learner's active vocabulary is a "good thing," and we have tried to simulate one obvious approach to achieve this aim. The simulations have suggested that our traditional way of thinking about this might not be taking into account some fundamental properties of networks, and that vocabulary networks might be more resistant to change than what these traditional approaches suggest. We have also begun to question what is really going on when students are asked to carry out traditional vocabulary exercises. More importantly, perhaps, we find ourselves being pushed in the direction of seeing a vocabulary network as a dynamic structure, one where a lot of interesting things are taking place at the edges of the network. We will take these ideas further in the next Part of the Workshop.

References

Meara, P. M. & Miralpeix, I. (2022). Vocabulary networks workshop 1: Introduction. *Vocabulary Learning and Instruction, 11*(1), 17–34. https://doi.org/10.7820/vli. v11.1.meara1

Brezina, V. & Gablasova, D. (2015). Is there a core general vocabulary? Introducing the New General Service List. *Applied Linguistics, 36*(1), 1–22. https://doi.org/10.1093/applin/amt018

Carter, R. (1987). Is there a core vocabulary? Some implications for language teaching. *Applied Linguistics, 8*(2), 178–193. https://doi.org/10.1093/applin/8.2.178

Foster, P. (2009). Lexical diversity and native-like selection: The bonus of studying abroad. In B. Richards, M. Daller, D. Malvern, P. Meara, J. Milton & J. Treffers-Daller (Eds.), *Vocabulary studies in first and second language acquisition* (pp. 91–106). Palgrave.

Gesa, F. (2019). *L1/L2 subtitled TV series and EFL learning: A study on vocabulary acquisition and content comprehension at different proficiency levels.* PhD thesis, University of Barcelona.

Meara, P. (1990). A note on passive vocabulary. *Second Language Research, 6*(2), 150–154. https://doi.org/10.1177/026765839000600204

Melka, F. J. (1997). Receptive vs productive aspects of vocabulary. In N. Schmitt & M. McCarthy (Eds.), *Vocabulary: Description, acquisition and pedagogy* (pp. 84–102). Cambridge University Press.

Milton, J. & Meara, P. (1995). How periods abroad affect vocabulary growth in a foreign language. *ITL Review of Applied Linguistics*, *107*/*108*, 17–34. https://doi.org/10.1075/itl.107-108.02mil

Miralpeix, I. (2020). L1 and L2 vocabulary size and growth. In S. Webb (Ed.), *The Routledge handbook of vocabulary studies* (pp. 189–206). Routledge.

Nation, I. S. P. (1990). *Teaching and learning vocabulary.* Newbury House.

Nation, I. S. P. (1996). *Vocabulary lists.* Victoria University of Wellington, English Language Institute, Occasional Publication n.17.

Pimsleur, P. (1967). A memory schedule. *The Modern Language Journal, 51*(2), 73–75.

Skehan, P. (2009). Lexical performance by native and non-native speakers on language-learning tasks. In B. Richards, M. Daller, D. Malvern, P. Meara, J. Milton & J. Treffers-Daller (Eds.), *Vocabulary studies in first and second language acquisition* (pp. 107–124). Palgrave.

Waring, R., & Takaki, M. (2003). At what rate do learners learn and retain new vocabulary from reading a graded reader? *Reading in a Foreign Language, 15*(2), 130–163.

Zaytseva, V., Pérez-Vidal, C. & Miralpeix, I. (2018). Vocabulary acquisition during study abroad. A comprehensive review of the research. In C. Sanz & A. Morales-Front (Eds.), *The Routledge handbook of study abroad research and practice* (pp. 210–224). Routledge.