Proceedings from the Second Annual Lucid Dreaming Symposium
Session 1: What is a Lucid Dream: Psychological and Physiological Considerations

EEG and Other Physiological Findings

Stephen LaBerge
Stanford University

Andrew Brylowski
University of Texas, Houston

LaBerge: The basic question we had in mind was, "What is happening in the brain when people become lucid?" We have had some ideas about this for many years. Earlier work that we had done indicated that lucid dreams are initiated during periods of cerebral (and autonomic) activation. It can be seen in a typical polygraph example of a lucid dream initiation that on the EOG channel, just a few seconds before the eye signal there is a striking suppression of finger pulse amplitude. It is like a switch is suddenly turned on--sympathetic activation. Likewise, you see a change in the respiration pattern; it changes from a regular rhythm and high amplitude to irregular and low amplitude with an increased rate. Also the top EEG channel will show a large skin potential artifact at the onset of lucidity. All of this indicates that when lucid dreams are initiated there is an activation of the brain. However all examples of the same process may not be identical. For instance the change in respiration may not be as clear.

Another indication is the H-reflex, which we have found is suppressed during the initiation of the lucid dream. It is more suppressed, statistically speaking, during the entire lucid dream than it is in REM sleep on the average. It is probably not more suppressed than it is in phasic REM, so basically we see that lucid dreams happen during periods of increased activation. There is greater eye-movement density, more autonomic variability; therefore, higher activation of the brain.

The data shown in Figure 1 represents standard scores averaged over 76 lucid dreams and 13 subjects. For each of these lucid dreams we drew a line at the onset of the lucid dream and divided the REM periods up into thirty second epochs before and after lucidity onset. The black line running through the center of each of these histograms indicates the onset of lucidity. Each histogram bar, then, is a thirty second average grand mean of standard scores. The top panel shows eye-movement density. You'll notice that there is a significant elevation of eye-movement density in the thirty seconds before the lucid dream starts. Likewise, there is a significant elevation of respiration rate; there is an elevation, but not significant, of heart rate, and there is skin potential activity. Also, you see in the first thirty seconds of the lucid dream, there's the same kind of activation, in fact, even a larger one. This activation is maintained at a higher level throughout the lucid dream. However, I wouldn't want to say that the high level of activation that you find later in the lucid dream is because of lucidity. We don't really know that; it could be due to dream content. But it seems clear that there is an association between lucidity and activation, especially if you look at eye-movement density, which may be the best measure of CNS activation up to this point. In something like 78 out of 80 of the lucid dreams, eye-movement density in the thirty seconds before the lucid dream was above median. What we conclude from this is that there is a necessary condition for people to realize they are dreaming, and that is sufficient activation of the brain. It seems that the lower levels of activity you get with, say, tonic REM, when there is not a lot of eye-movement activity, is not sufficient for people to become reflectively conscious.

This is all background information, and what we have known for several years about the initiation of lucidity. Clearly, the brain must be activated, and especially it must be activated in the first thirty seconds of the lucid dream and the thirty seconds right before it. Something is obviously happening in the brain and we would assume that the cortex would show it. But we had little idea, other than from this autonomic data, about what exactly was taking place, whether the entire brain was equally activated or whether some specific areas were more activated than others.

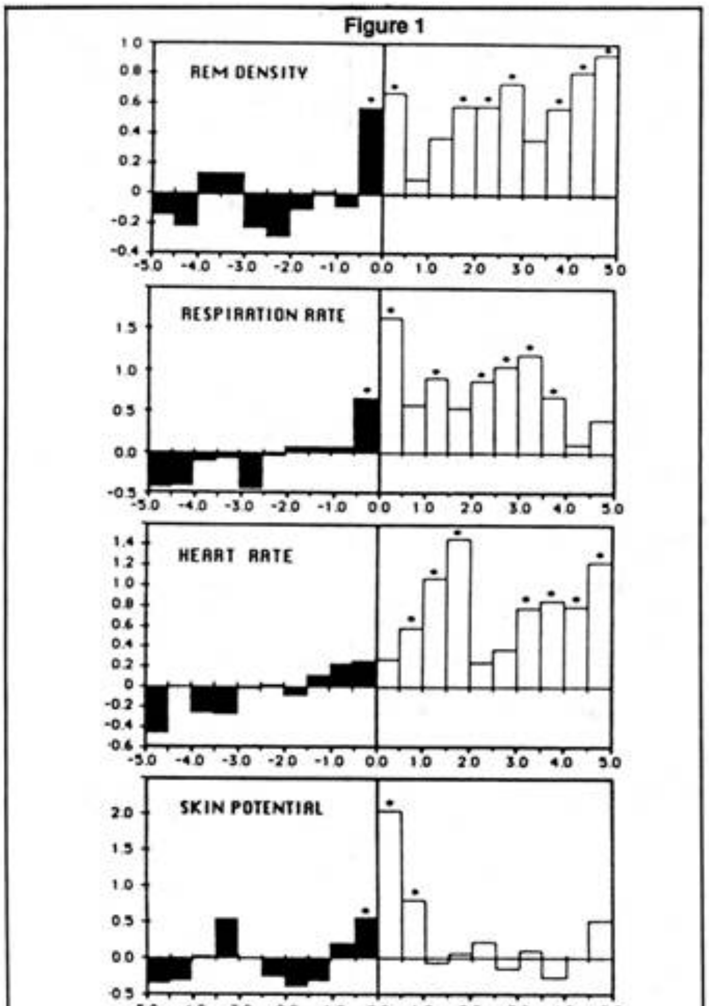

Figure 1 Histograms of mean Z-scores for EM, RR, HR, and SP. Bins are 30 seconds in length with t=0 representing the signaled onset of lucidity. Ns vary with variable and bin, but all values are averaged across lucid dreams and subjects. (*p<.05)

There is a use for this information beyond the obvious research interest. Since we know that lucid dreaming never takes place unless there is a sufficient level of eye-movement activity, for example, we know that when we're trying to induce lucid dreams by external stimulation that we don't want to apply a stimulus during a period of low eye-movement density. But we could probably improve on our induction technique if we were to find that there were certain areas, that is if the left frontal cortex, or the left temporal lobe, or some other area was reliably activated in the thirty seconds before a lucid dream emerged. We could then watch for when that activation occurred and apply a stimulus, or a reminder just at that point and perhaps get much better results than we are currently.

So, we were interested in finding out which parts of the brain would be more activated in lucid dreams. Technology has developed in the last several years to the point that it is possible to do extensive cortical brain mapping. You can now collect data from, say, twenty-eight, or thirty electrodes on the cortex, and then make maps showing the distribution of the brainwave activity in various states. This is what we have done. Dr. Brylowski is going to describe to you how the data was collected.

Brylowski: One of the considerations in doing this study was the time it would take to apply the electrodes and collect the data. Further we felt it was necessary to do this consistently over a period of time. That is to have the electrodes applied to the same place consistently over time and to be able to do it in the future in a relatively rapid manner so that any people who would participate could be monitored with relatively little inconvenience.

One of the more recent technologies relevant to these concerns is the electrode cap, which is basically like an expandex kind of hat with electrodes in it. One of the criticisms of this hat is that it slides around and is not a very good tool except for short term recordings, of perhaps five or ten minutes. So one of the initial technical obstacles was to be able to put this on and record and have impedance, or resistance at the electrode, which was consistent and low through an eight hour recording of a nights sleep.

We found that this cap could be adapted by taking the little styrofoam doughnut that came with this cap for absorbing sweat on the forehead, and putting them all over the cap. Further we took regular electrode cream and put it on the electrodes around the periphery and around the center where the wires actually come off the cap so that when the subject roles around in bed or tosses and turns the electrodes are not displaced and a quality recording can be obtained for the whole night.

The other problem with this type of technology is that it’s very difficult to calibrate the actual spot on the scalp where the electrode is, beforehand. There is a product called omniprep, which electroencephlographic technicians like to use, which helped solve this problem. A little wooden Q-tip stick that could be dipped in the omniprep beforehand fits perfectly in the little holes of this cap. So this cap could be applied and all twenty eight electrodes put on the scalp in about twenty minutes.

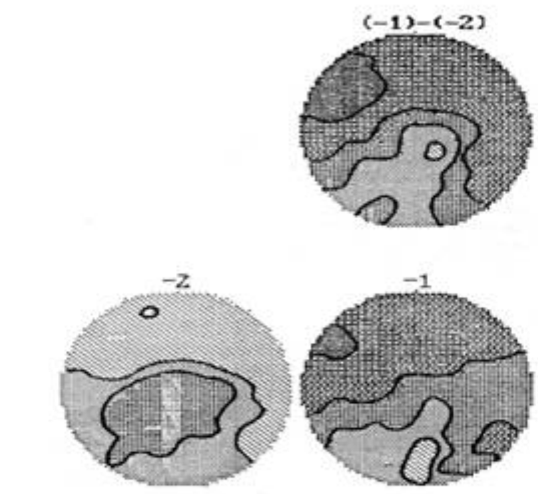

Figure 2A

The technique is basically very simple. Once the scalp is abraided and is ready to be injected with electrode paste the rest of the electrodes can be put on very quickly. The impedance can be checked. Most of the electrodes were either four ohms or lower in impedance and that was maintained throughout the night and in the morning.

Then the subject could easily go to bed and be hooked up. This whole process would take between twenty and forty minutes. I'm sure it could be done even quicker with added practice. The subject was a twenty eight year old fourth year medical student who was a frequent lucid dreamer and had very good dream recall. I'll turn it back over to Stephen who will interpret some of the results.

LaBerge: Yes, well, you can get medical students to do almost anything, you see. Now, Figure 2 displays some of the data we have analyzed. What we did corresponds to the histogram of Figure 1 which you saw earlier. The records were divided at the point of the start of the lucid dream, and this shows four thirty second periods, two before and two after the moment of initiation. The first thing we did was average together across lucid dreams all of the two and a half second intervals of digitized data, the 256 point FFT's, and then we computed standard scores. This maps we are seeing are of standard scores, so you see, if you look at the grey scale on Figure 2, the values for all of the shades displayed in the maps add up to zero, which is represented by the black.

We looked at five different frequency bands of EEG. You can divide any complex wave form into a set of simpler waves, and brain waves are usually divided up into certain frequency bands, which are the ones we used in this analysis. Our results for some of the bands were hard to interpret, mostly due to the presence of eye-movement artifact in the frontal areas. One of the next things we're going to have to do, of course, is to separate out the eye-movement activity in order to see whether there is anything else going on in these bands besides that. Unfortunately, the particular commercial machine that we were using did not allow us to save the raw data or the EOG or eye-movement activity, so we couldn't do that in this case.

The bottom four maps, labeled -2, -1, +1, and +2 are 30 second epochs with +1 being the first 30 seconds of the lucid dreams as marked by the signals the top three maps on this figure show t-scores of the difference between the minus two condition and each of the other conditions. The N's in this case are five for each map because we have combined together five lucid dreams. Once we had standardized them, transformed the data into z-scores, we averaged the dreams together. The minus two condition is our control condition, and is supposed to represent a random selection of REM sleep not necessarily associated with a lucid dream.

If you look at the histograms in Figure 1 you'll notice that if you go back to minus two minutes or so before the lucid dream you are at the mean level of activation. There is no reliable significant increase or decrease of activity at that point. This is what we are using for comparison in the EEG analysis. One limitation of this study is that in the case of the autonomic analysis we had averaged entire REM periods, or at least thirty minutes of REM before the lucid dream, if there was more than that. So, we had a more reliable estimate of the overall mean for the entire REM period in that study than we do in this case, where we may have only something like ten or fifteen 256 point FFT's averaged to make our reference. Therefore, there is a bit of added variance in our comparisons here, which simply makes the probabilities a bit larger than we'd find with a better average. The particular patterns seen would be the same.

The picture you are seeing in Figure 2 did not come from the NeuroScience Brain-Mapper that the data was collected with. Instead, this is a program written by Romana Machado in my laboratory for an IBM AT for plotting the data that was rather laboriously copied from a print-out from the NeuroScience machine by Lynne Levitan, who typed approximately twenty-five thousand numbers into a computer, and then we read it again. This was just in order to get this analysis done this week so that we could show it to you. It's not the best way to transfer data!

Now, let's look at the most interesting frequency band--alpha, or 8 to 12 Hertz. Alpha seems to be one of the few sets of brainwaves for which there is some consensus about it's interpretation. Its presence is generally interpreted to tell something about the activation of the brain--that when there is less alpha in a region, the brain is more activated in that particular region. Because of this inverse relationship between alpha power and brain activation, we have the scale reversed on this figure. At the top of the scale we have black, representing higher levels of alpha, and at the bottom is white, again representing higher levels of activation and, in this case, lower levels of alpha.

Now, there is some relative activation apparent in the minus two frame, which I think is due to the fact that the sample is poorer than we'd like. It actually has more variance than the other frames. The variance is smaller for frames minus one and plus one, because the activation is reliably higher at the onset of lucidity. In the earlier frame, minus two, there's a lot of random variation. Anyway, you'll see that in the first frame of the lucid dream there's a significant elevation of activation in the left hemisphere, especially in the left parietal and posterior temporal cortex. This is consistent with what we found with a few other lucid dreams we had analyzed at Stanford, where, essentially, the only difference was in the ratio of alpha power in the left versus right parietal lobes, just like this data shows.

I'm not going to tell you what this left-parietal activation means now, for two reasons. One is that I hope our commentators might have something to say about what it might mean, and two is that there is one other problem with this data, in terms of the spatial distributions. Now, what we're seeing here does not tell us what the absolute levels of any of these powers are at any particular electrode sight because they have been converted to z-scores before averaging. At any one point, we are only looking at the difference in power. So, it's a little hard to say if this is relative left parietal activation or if left parietal activity has actually been low and is simply coming up to a normal level at the beginning of lucidity. We cannot tell which is the case from this data. A further problem is that this particular brain mapping machine uses linked ears as a reference, causing distortions in the field, and making it difficult to determine absolute levels. This is a direction for future research.

We also looked at the total power of all the frequencies excluding the delta band. I think this is also contaminated with eye-movement activity, because it includes theta power. We found apparent activity in the right frontal cortex which is probably eye-movement artifact. It is interesting, however, that it was found on only one side of the brain. One thing that has not been looked into is whether or not there is a difference in eye-movements to the right or left at the initiation of lucid dreaming. We would expect that lucid dreaming is a primarily left hemisphere process, that it seems to require more left hemisphere style cognition than ordinary dreaming. To attain full lucidity, you have to spell out to yourself, "This is a dream," which is a specifically linguistic, and therefore, left hemisphere, task.

Question: Was there any specific activity that the lucid dreamer was doing after he became lucid in these particular dreams?

Brylowski: That's a very interesting question. One of the reasons I had for doing this analysis was to move towards developing a model of lucid dreaming by which we may be able to predict where lucidity occurs without having to mark lucidity onset with, for example, eye-movement signals. Another reason I had was to see if a task performed in waking, such as fist clenching or moving a finger, would produce similar EEG activation as the same task performed in the dream state. That data hasn't been analyzed yet.

Most of these maps, since they cover the first minute after the initiation of lucidity, did not contain any specific task beyond the spontaneous dream activities occurring at that time. In some of the dreams, after a period of lucidity, I did perform certain tasks or pre-planned dream-imaged activity, but they do not appear in these maps. Again, this is an area of future research, to look for any differences between the grand averaged maps of lucidity initiation and the various types of tasks that could be performed in the dream state.

Question: Could you explain in more detail exactly how you are ascertaining when the lucid dream begins?

LaBerge: We determined where the lucid dream began primarily from an eye-movement signal: the left-right-left-right signals. In some cases, if a few seconds before the signal there was a very clear autonomic activation that accompanies a person realizing, "Aha! This is a dream!", we drew the line there instead of at the eye-movement signal. If there wasn't a better indication than the eye-movement signal, then that is what we used.

Question: First, when in the night do lucid dreams occur? Do they occur in all REM periods or do they concentrate in the morning? And, are there changes in the sleep architecture with lucid dreaming? For example, are there changes in stages three and four, or changes in total sleep time, or total REM time?

LaBerge: To take the last question first, changes in architecture: we looked into sleep architecture of this same sample of eight lucid dreams we have described here, and found no particular differences in stage two or three sleep. The major distribution of lucid dreams is towards the end of the night, but you find lucid dreams in all REM periods. I have a paper in which I describe finding a linear relationship between the REM period number and the likelihood of lucid dreams. This is probably related to a circadian rhythm, the fact that the REM activation cycle reaches a maximum at about 10 to 11 hours after sleep onset, meaning that as the night goes on lucid dreaming becomes easier. The results I have just mentioned are, of course, confounded by the fact that they are not really based on REM period number, and may be strongly affected by this time of night effect.

Hunt: Jayne and I were just discussing two points regarding your findings. One is that some of your evidence of dual-sided parietal activation may be consistent with findings of increased kinesthetic sensation in lucidity in general. The other is that I think Bob Ogilvie and Paul Tyson will be very pleased because together we found, with more standard filtering for alpha, an alpha effect. Although we had very few really good lucid dreamers, and most of our episodes were pre-lucid, meaning questioning reality in the dream, but not being certain, we found a similar enhancement of alpha.

LaBerge: Well, actually, if we return to the slide on alpha activity (see middle map of top row of Figure 2), what we're seeing is less alpha, a decrease of alpha at the initiation of lucid dreaming. It's a different question whether there's a global difference in alpha power in a REM period that has a lucid dream in it. Your measures were of global alpha power, but these are a measurement of relative differences in the amount of alpha at different periods. What we have found is a decrease in the relative amount of alpha at the initiation of the lucid dream, whereas what your studies showed was that if you looked at the magnitude of alpha in the lucid or pre-lucid REM period, not the relative change at initiation, it was higher than for non-lucid REM periods.

Question: I want to ask you how the lucidity was initiated. If it was initiated from waking, clearly this would have made a big difference in terms of alpha. My second question is about the subject signaling lucidity with eye-movements. How do you think the data would have been different if he were to have signaled in some other way, such as by a change in respiration rate, or a clenching of muscles in the hand?

LaBerge: In answer to your first question, this sample did not include any lucid dreams initiated from waking. As for your second question, I don't think the eye-movement signals should make that much difference for the alpha maps, because most of the eye-movement artifact occurs in the delta band. That, of course, seriously confounds the delta maps and to a certain extent the theta maps. In future work we plan to have our computer system automatically compensate for eye-movement artifact.