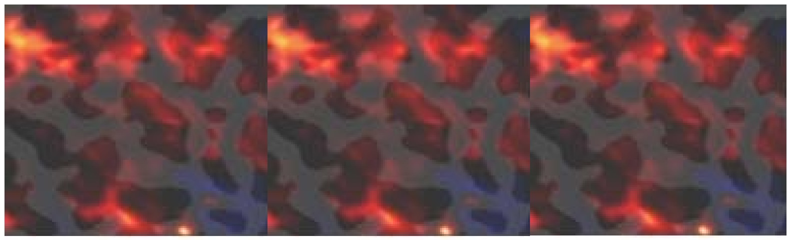

Cognitive Science Online

Vol.1, Issue 2
http://cogsci-online.ucsd.edu

Letters

Articles

Information

Cognitive Science Online is an online journal published by the UCSD Cognitive Science Department and seeks to provide a medium for the cognitive science community in which to exchange ideas, theories, information, advice and current research findings. This online publication is a peer-reviewed and highly interdisciplinary academic journal seeking contributions from all disciplines and methodologies investigating the mind, cognition and their manifestation in living, and possibly artificial, systems. For more information about this journal, submissions, back issues, please visit our website at http://cogsci-online.ucsd.edu

Contact Information

Department of Cognitive Science
University of California San Diego
9500 Gilman Drive
La Jolla, CA 92093-0515
cogsci-online@cogsci.ucsd.edu

Editors

Advisory Editorial Board

Letter from the Editors

Our civilization, our so-called culture, is vulnerable, forming but a thin veneer over the chaos which waits unsleeping and unblinking beneath, watching for the opportunity to escape. The ancients knew this well. Perhaps because they lacked a concept of linear progress they felt nearer to their origins which they saw all about them. For those cultures, chaos was ever present and had to be constantly combated. The ancient Babylonians re-enacted the defeat of chaos each year during their great Spring Festival. This reminded every citizen of what had been necessary to build civilization. The ancient Egyptians saw chaos as the realm of Seth, and civilization was ensured by Ma'at, goddess of balance and harmony, truth and justice.

Today, however, we often feel confident that we are at the end of a long path of philosophical and technological progress, and that the underlying chaos has been suppressed, arising only in mythologies and storytelling as demons, monsters, or "the other," which must be vanquished by our heroes. The truth behind the metaphors has been forgotten however, and we often appear to think our civilization so strong that we can treat it with disdain. The events in the Balkans during the course of the 1990s showed the foolishness and vacuity of this approach. It showed just how thin was the cover over chaos and destruction.

Which brings us to journalism, for it is, above all, a civilizing tradition. It remains beyond the divisive and contrary forces of politics and religion; beyond the relentless currents generated by the ambition of kings, presidents and generals; beyond the confident rigidity of inflexible and self-important belief systems. Journalism trusts in the person. It celebrates the innate dignity of every human being. A dignity which needs not just recognition and tolerance from others but a physical, intellectual and spiritual space in which to grow. For the Journey of a journalist is a growth towards awareness and insight.

But something weird seems to be happening around us. There are wars and rumors of wars, chaos and rumors of chaos. Who knows where it will all end? In our confusion - or our certainty - we are no different from many people in many ages past. What is important is the maintenance of our precious civilization, which has a significance beyond national and ideological boundaries. What is important are those landmarks which keep us civilized - morality, responsibility and information - the very concerns of journalism.

Since our inaugural issue, Cognitive Science Online has benefited from the welcome additions of David Groppe and Arielle Borovsky, and so our editorial staff has nearly doubled in size to five stellar graduate students, venturing forth to fight chaos, armed with enthusiasm, integrity and razor-sharp repartee. The journal has grown precociously from the mere sapling concept it was a year ago into the tree of knowledge that it is today, shooting up through the weeds of destructive chaos, nurtured only by the blood, sweat and tears of its contributors as it grows towards the sun of awareness and insight. What you see before you is the culmination of our efforts to keep chaos at bay, to bring you pure, unadulterated information, so that you too may climb aboard this tree, fending off the dark tendrils of chaos that might otherwise pull you down into the mire.

However, magnificent as our efforts are, we, the editors, still rely on you, the readers, to provide us with the raw material to sustain us all in our struggle against chaos. We need your submissions or our journal will wither away and die, and chaos will once again reign supreme. Send us any material you would like to see published in this medium. Send us your papers, editorials, humor, fiction, questions, suggestions, comments, etc. The viability of our journal depends on the nourishment it receives from you, our readers, as we editors are merely the humble shepherds of information, here to serve you and the entire cognitive science community by caring for this tree during its long journey towards the firmament.

Christopher Lovett

Ayşe Pınar Saygın

Hsin-Hao Yu

Arielle Borovsky

David Groppe

Cognitive Science Online, Vol.1, pp.34–45, 2003 http://cogsci-online.ucsd.edu

Blending in action: Diagrams reveal conceptual integration in routine activity

Beate Schwichtenberg
Department of Cognitive Science
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0515
bschwich@cogsci.ucsd.edu

1 Introduction

Picture a first-time participant in a "brainwave" study. He is seated in the recording room, the experimenter shows him his brainwaves, and then leaves the room. The student watches his EEG with fascination until abruptly all the lines on the screen go flat. Does this mean his brain stopped working? His confidence only returns when the experimenter reassures him that she blocked the transmission of his brainwaves onto the screen.

Why was the student afraid of being brain dead? The student did not distinguish between his brain activity and the lines on the computer screen. He treated brain activity and lines on the screen as one unique thing, and thus inferred that flat lines must mean no brain activity. Conceptual integration, or blending, is a framework for the analysis of phenomena such as this, where information from two separate domains is brought together and integrated, producing emergent structure and generating new insight (Coulson & Oakley, 2000; Fauconnier & Turner, 2002).

In the present paper, I will discuss conceptual integration processes in everyday activity. The examples come from a routine activity in a cognitive neuroscience laboratory, preparing a participant for the recording of his EEG. The experimenter needs to work with several artifacts, such as an electrode cap and an impedance meter. The coordination of these artifacts is made easier by particular diagrams and charts located throughout the lab. My central claim is that activity and conceptual integration mutually influence each other. On the one hand, the intensity and extent of interaction with and coordination between artifacts finds a reflection in the extent of integration between the two domains. Activity gives rise to blends. On the other, the integration of domains may lead to action that would not have been performed in either of the original domains. Blending gives rise to activity.

2 Background: The Capping Process

In order to understand the discussion that follows, a short and partial overview of the artifacts used and actions taken during the process of preparing a participant for the recording of the EEG is needed. Two caveats up front: First, this description is not intended as an accurate description of technical details or guidelines for the recording of the EEG. Technical descriptions, if included, are often simplified. They aim to provide insight on the more informal understanding that an experimenter might draw upon during the setup process. Second, the artifacts and processes may not be the same in all labs using the methodology.

Figure 1: Electrode Cap with Quick Inserts

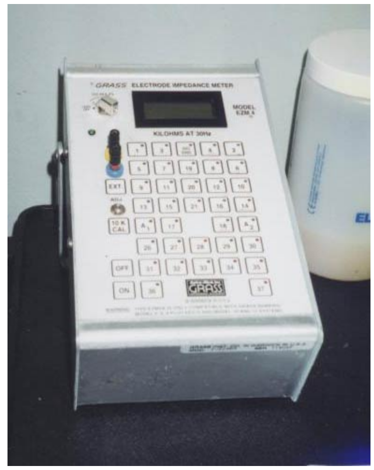

Figure 2: Impedance Meter

This includes, among others, the diagrams and charts in the lab as well as the number and location of the electrodes on the cap. This paper is not intended to make any general claims about the recording of the EEG. Instead, it aims to show that blending and activity can interact in routine activities.

Preparing a participant for an experiment is usually referred to as *capping*, because the main task is putting an electrode cap onto the head of the participant. The *electrode cap* (Figure 1) is a cap with several electrodes that are spatially arranged in a special configuration. It resembles a tight-fitting swimming cap with chinstraps and little cylinders containing the electrodes. The electrodes, which are arranged on the cap in a geometric pattern resembling four concentric circles, are placed within white plastic cylinders. Each cylinder has a small whole in the middle that will be filled with conducting gel in preparation for the experiment. The wires of the electrodes are threaded through the cap, and combined in two flat band cables. Each cable combines half the electrodes of the cap.

The *impedance meter* (Figure 2) is used for checking the impedance of the electrodes. It resembles a bulky pocket calculator. It has about 40 buttons, which correspond to individual electrodes in the cap. A panel near the top of the box displays the impedance at the electrode when the corresponding button is pressed.

The *numbered layout* (Figure 3) is a diagram that maps the correspondence between the electrodes and the impedance meter. The diagram is a schematic top-view drawing of am electrode cap as it sits on the head. The nose and the ears are sketched on the top and the left and right sides, respectively. The diagram contains four concentric circles, mirroring the layout of the electrodes on the cap. The electrode locations are noted as little circles embedded in the concentric circles at the appropriate location. Each circle contains the number of the corresponding button on the impedance meter.

Recording good data requires that the impedances of the electrodes are below a certain threshold. *Lowering impedances* is one of the main tasks during the capping procedure. Three routine steps for lowering the impedance of an electrode are *moving the hair* below an electrode to the side, *squirting conductance gel* into the electrode with a syringe, and gently *scratching* the skin below the electrode with a sharp needle. The *impedances* are *checked* using the impedance meter. If the impedance (displayed on the panel of the impedance meter) is above threshold, the standard procedure is to continue scratching the skin below the electrode and/or to apply more gel to the electrode. Sometimes, however, an electrode

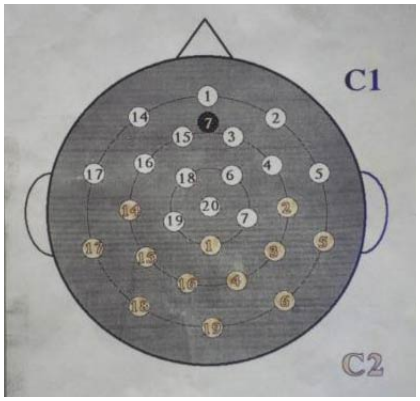

Figure 3: Numbered Layout

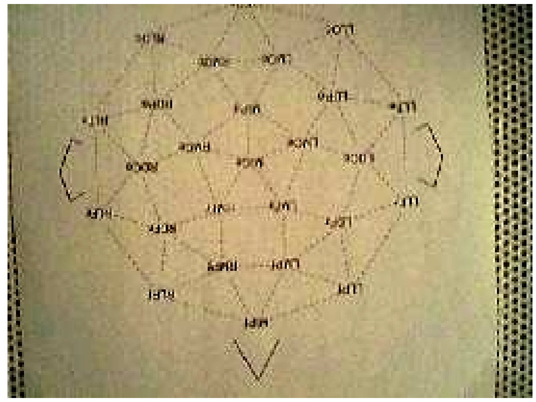

Figure 4: Labeled Layout

may be faulty, rendering it impossible to lower the impedance below threshold. In this situation, a *quick insert* is used, which overrides the cap electrode. The quick insert is inserted in the cylinder containing the electrode. It is usually fastened to the cap with a strip of tape. Figure 1 features a quick insert in one of the electrodes.

Capping is done in a small room adjacent to the recording chamber in which the experiment itself is conducted. The participant, wearing the electrode cap, sits on a chair. The numbered layout is taped to one of the walls. The impedance meter is easily portable, thus does not have a standard location. The process of reducing all impedances below threshold takes half an hour or longer. During this time, the electrodes on the cap and the impedance meter need to be constantly coordinated. The numbered layout displays the (nontrivial) correspondences between the impedance meter buttons and the electrode locations.

After all impedances are below threshold, the electrodes need to be connected to the EEG recording equipment. This takes place in the recording room itself, where the participant is seated in a comfortable chair, facing a computer monitor. Again, several different artifacts and diagrams are used in the process.

The *labeled layout* (Figure 4) is a diagram similar to the numbered layout. Instead of numbers, it gives names for the individual electrodes. The names are associated with the electrodes throughout the preparation and later the data analysis process. The electrode labels roughly correspond to the parts of the brain over which the electrode is located. For instance, the electrode placed over the left outer part of the prefrontal lobe is called LLPF, short for left lateral prefrontal.

The *connector boxes* (Figure 5) are two square boxes, located in the back of the recording room. Each connector box contains plugs in a rectangular grid pattern. Each connector has a number, and a wire sticking out of it. Connectors on a box can be uniquely identified either by their number or by the unique wire color/connector color combination. The wires from each connector box are combined in a flat-band cable. For the EEG recording, these cables are connected with the cables from the cap.

The chart labeled *connectors for cap/quick inserts* (Figure 6) is used to cross-reference the connector boxes and the cap electrodes. This mapping is needed whenever a quick insert is used. The electrodes are labeled according to their names described in the labeled layout. If there is a quick insert in one of the electrodes, the connection of the cap electrode to the connector box needs to be replaced by the wire coming from the quick insert. The cross-reference table gives the connector number corresponding to the electrode label.

The labeled layout, the connector box, and the connectors for cap/quick inserts chart are

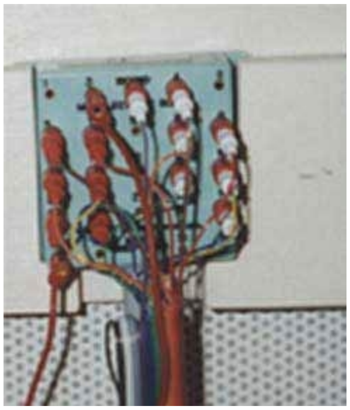

Figure 5: Connector Box

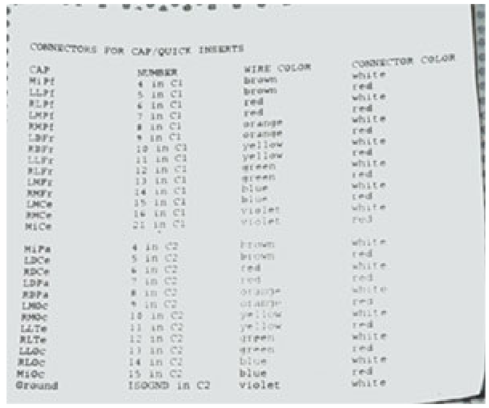

CONNECTORS FOR CAP/QUICK INSERTS

CAP	NUMBER	WIRE COLOR	CONNECTOR COLOR
MiPf	4 in C1	brown	white
LLPf	5 in C1	brown	red
RLPf	6 in C1	red	white
LMPf	7 in C1	red	red
RMPf	8 in C1	orange	white
LDFr	9 in C1	orange	red
RDFr	10 in C1	yellow	white
LLFr	11 in C1	yellow	red
RLFr	12 in C1	green	white
LMFr	13 in C1	green	red
RMFr	14 in C1	blue	white
LMCe	15 in C1	blue	red
RMCe	16 in C1	violet	white
MiCe	21 in C1	violet	red
MiPa	4 in C2	brown	white
LDCe	5 in C2	brown	red
RDCe	6 in C2	red	white
LDPa	7 in C2	red	red
RDPa	8 in C2	orange	white
LMOc	9 in C2	orange	red
RMOc	10 in C2	yellow	white
LLTe	11 in C2	yellow	red
RLTe	12 in C2	green	white
LLOc	13 in C2	green	red
RLOc	14 in C2	blue	white
MiOc	15 in C2	blue	red
Ground	ISOGND in C2	violet	white

Figure 6: Connectors for Cap/Quick Inserts

Figure 7: Brainwaves. Resetting MiPF.

mounted onto the wall behind the participant's seat in the recording chamber. The labeled layout is taped to the wall above the head of the participant. The upside down orientation of the labeled layout helps identify electrodes at a glance: this way, the right side of the participant corresponds with the electrodes labeled right in the diagram.

After the scalp electrodes are connected to the recording equipment, the EEG needs to be processed, displayed and recorded on a computer. These preparation processes are invisible to the participant: the amplifiers and computers are located in the room adjacent to the recording chamber. The experimenter can chose to display the EEG on the computer screen which the participant is facing.

The *amplifier* is part of the transmission of brain activity onto the computer screen. Brain activity is measured as voltage changes over time. Each *channel* corresponds to the voltage difference between one of the electrodes and a reference electrode, and this difference is eventually displayed on the screen. The amplifier contains a *reset button* for each channel, which prohibits the signal from passing through the amplifier.

A *computer screen* (Figure 7) is used to monitor the EEG during the recording period. On the screen, each channel is labeled with the name of the electrode whose signal it represents. The EEG itself is represented as wiggly lines on the screen. Pressing the reset button on the amplifier flattens the line on the screen. Figure 7 shows this for the electrode labeled MiPF.

3 Conceptual Blending: The Brainwaves Blend

Conceptual integration (or conceptual blending) is a framework for the analysis of higher cognitive phenomena, in which selected parts of two or more conceptual domains are brought together and combined ("blended"), producing emergent structure and allowing insight that is not contained in either of the original domains. Conceptual blends can be represented with conceptual integration networks. These networks contain two or more input spaces representing the different conceptual domains. Corresponding elements in the input spaces are connected via *cross-space mappings*. The blended space represents the integration of these spaces. Finally, the generic space contains structure that is shared among the input spaces and the blend.

I will use the Brainwaves blend as a case study of the blending process. First, I will discuss the Brain Electrodes blend, which serves as one of the input spaces for the Brainwaves blend. Next, the Brainwaves blend proper. Third, I will show how this blend can give rise to actions that are not meaningful in either of the input spaces alone.

The Brain Electrodes blend captures the idea that scalp electrodes register brain activity. The *input spaces* in the Brain Electrodes blend are the Brain and the Electrode Cap. The conceptual domain Brain provides elements such as the location of neural tissue, the activity of cortical neurons, and the changes in the electromagnetic fields that are caused by neuronal activity. The Electrode Cap provides elements such as the spatial layout of the electrodes in the cap (and thus on the scalp), and the signal that the electrodes record. The Brain and the Electrode Cap are related on many levels. For instance, the electrodes are placed on the scalp above certain standardized locations in the brain. Further, the signal that the scalp electrodes register is partially caused by the activity of cortical neurons. The correspondences between these domains are realized as cross-space mappings in a conceptual integration network. For instance, brain activity and the signal reading at the electrode are linked by a cause-and-effect cross-space mapping. The conceptual integration network for the Brain Electrodes blend is depicted in Figure 8. The two input spaces to this blend, Brain and Electrode Cap, are represented as circles. The cross-space mappings between the elements in these two domains are represented by lines, connecting the elements. Of course, there are elements in each input space that do not have a correspondence in the other input space: the colored fabric and the chinstrap of the Electrode Cap does not correspond to anything in the Brain input.

The *blended space* represents the integration of the Brain with the Electrode Cap. In the blend we can say that the electrodes register brain activity. The signal at the electrodes is identified with brain activity. Thus, neuronal activity and the readings at the electrodes, which were linked by a cross-space mapping in the input spaces, are *compressed into uniqueness* in the blend. That means that there is no distinction between brain activity and the signal at the electrode. Further, in the blend the electrodes are located above certain brain areas. The electrodes are not seen as placed on the forehead or over the back of the head, but as placed above the frontal or the occipital lobe. The blended space is depicted as the bottom circle in the conceptual integration network (Figure 8). The lines between the input spaces and the blended space represent selective projections from each of the input spaces into the blended space. Both the compression of neural activity and electrode recording and the correspondence between electrode location and neural tissue are marked in the blended space.

The final element of a conceptual integration network is the *generic space*, which captures commonality between the input spaces. In the Brain Electrodes blend, the structure common to both input domains, Brain and Electrode Cap, is the existence of an event, the neuronal activity or the electromagnetic signal, at a specific location, the neural tissue or the electrode location. In the network, the generic space is represented by the top circle.

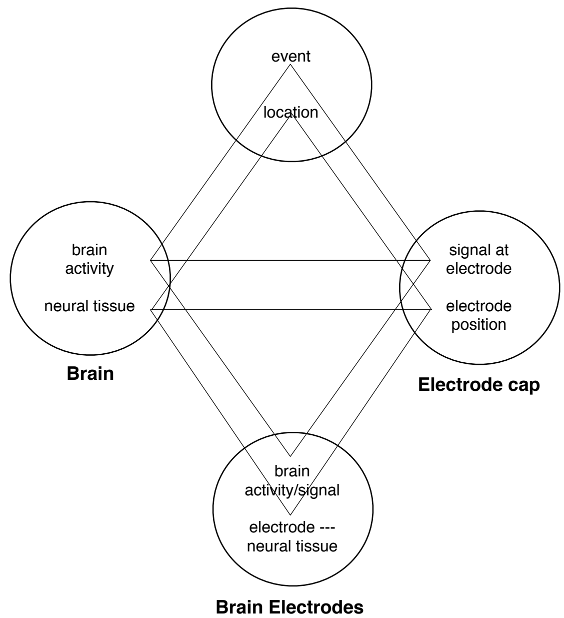

Figure 8: Brain Electrodes Blend

The generic space contains the elements and structures common to the input spaces, and therefore represents the most general underlying structures. As this general knowledge is not essential to any of the analyses below, I have omitted the generic space from my diagrams.

The Brain Electrodes blend is crystallized in the labeled layout (Figure 4). The diagram places labels, such as "LLPF" for left lateral prefrontal cortex, in a grid of concentric circles representing the positions of the electrodes on the cap. The two domains Electrode Cap and Brain are spatially superimposed, and thus spatially integrated. The labeled layout provides a stable, external representation of the correspondences between the electrode locations and the underlying brain areas. It serves as a *material anchor* (Hutchins, 2002), which is a physical object with spatial structure that brings stability into a blend.

The Brainwaves blend (Figure 9) captures the identification of the ups and downs of lines on a computer screen with the brain activity that is measured by the electrodes. The inputs in this blend are the Brain Electrodes, which register brain activity above specified parts of the brain, and the Screen, which features "wiggly" lines with labels. Changes in brain activity correspond to patterns of ups and downs of the line; the placement of the electrode (above a certain area of the brain) corresponds to the label to the left of the line. In the blended space, the inputs Brain Electrodes and Screen are combined. The patterns on the screen and the brain activity recorded by the scalp electrodes are compressed into uniqueness; we see the Brainwaves on the computer screen.

Consider the participant who was afraid of being brain dead when he saw his "brainwaves" go flat on the screen. Neither lines on a computer screen nor voltage changes by themselves could have triggered this reaction. Only in the Brainwaves blend, where the lines on the screen are identical to and inseparable from the activity of his brain, does this reaction occur. The student used the blend for an unpleasant, but creative inference that is not suggested by either of the input domains alone.

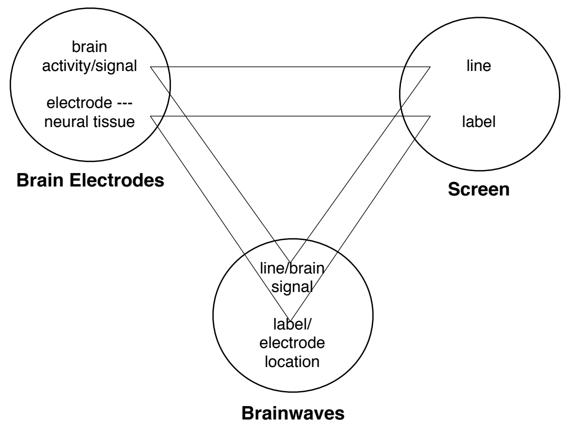

Figure 9: Brainwaves Blend

Emergent behavior that is not meaningful outside the Brainwaves blend can be observed with the experimenter, too, who has a much more sophisticated understanding of the concepts, for instance when dealing with electrode drift. Drift is a problem with registering the brain activity during the recording of the EEG. The problem originates directly at the electrodes. On the computer screen, drift is easily detectable. The lines on the screen, which are usually centered around a zero mark, quickly drift back and forth from the top to the bottom of the window and lose their zero-centering. Pressing the reset button on the amplifier flattens the lines on the screen: while the button is pressed, the line rests at zero. I have observed experimenters repeatedly pressing the reset button when they observe drift, in order to solve this problem.

Can pressing the reset button remove electrode drift? The problem occurs at the electrode, not on the screen. Recall that brain activity is transmitted in a unidirectional manner from the electrodes onto the screen. The signal is registered at the electrodes, changed and enlarged in the amplifier, and then displayed on the screen. Actions taken at later stages in this signal transmission can't change the signal at earlier stages. Pressing the reset button on the amplifier flattens the corresponding line on the computer screen, but it does not influence the way in which brain activity is registered at the electrodes. The reset button can not remove the conditions leading to drift, and thus is not a solution for the problem.

Why then does an experimenter press the reset button, if it does not solve the problem? Recall the Brainwaves blend: The key element in this blend was that the computer screen displays brain activity. Thus looking at the screen means looking at the brain activity, and looking at whatever the electrodes register. No distinction is made between the represented and the representation, they are compressed into uniqueness. Resetting a drifting electrode is emergent behavior arising from this compression. Resetting the channel zero-centers and flattens the line; the problem temporarily disappears. The compression of the line and the signal at the electrode leads to the inference that whatever causes the desired change to the line, causes the same change to the electrode.

In the following analysis (see Figure 10), the Amplifier is added as a new input domain to the Brainwaves blend. The channel numbers on the amplifier, the electrodes and their labels correspond to each other. The status of the reset button is a bit more complicated. Pressing the reset button interrupts the transfer of the electrode signal onto the screen. The screen then no longer displays brain activity, but rather a flat, zero-centered line. Pressing the reset button thus causes a change in the line. This causal relationship is captured in a

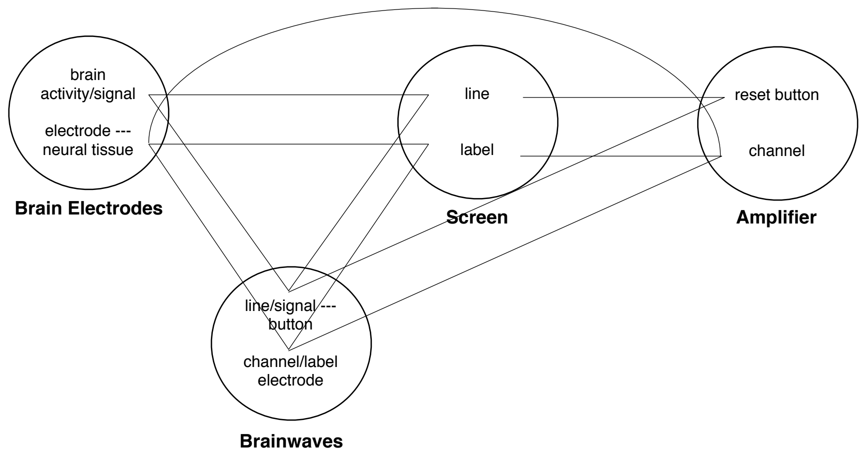

Figure 10: Backwards Causality

cause-effect cross space mapping. Note that the reset button is not connected to the signal at the electrode: reset button and brain signal don't influence each other. In the blended space, the electrode, its label and the amplifier channel are not distinguished: they are compressed into uniqueness. The electrode signal and the line on the screen are integrated in the same way: they are compressed into uniqueness. The reset button and its power to change the line on the screen are projected into the blended space. But note that in the blend the line on the screen is compressed into uniqueness with the brain activity measured at the electrode. The reset button causes a change with the line, but the line is not only a line: it is the measurement of brain activity at the electrode. Only in the blend does the reset button have an effect on the signal at the electrode: only in the blend can it be used to make the signal better. In the input spaces, the reset button does not and can not have this influence. It is at a much later stage in the unidirectional signal transmission process than the electrode itself.

The conceptual integration network is displayed in Figure 10. The input domains Brain Electrodes and Screen are the same as in the Brainwaves blend, the Amplifier is a third input domain. The electrode labels (in the Brain Electrodes and the Screen inputs) and the corresponding channel number on the amplifier are linked via cross-space mappings, indicating that they represent each other. Labels and channel numbers are projected into the blended space. They are compressed into uniqueness. As discussed in the Brainwaves blend, the signal at the electrode and its representation on the screen are linked by cross-space mappings and compressed into uniqueness in the blend. The line on the Screen and the reset button on the Amplifier are connected by a cause-and-effect relation. Pressing the reset button causes the line on the screen go flat. This relationship is projected into the blend: the line and the reset button in the blend are still causally linked. Note that the line on the screen and the signal from the electrode are compressed into uniqueness in the blend. There is no distinction between the electrode signal and its representation on the screen. Thus, the reset button is not only causally linked to the line, it is causally linked to the electrode signal which is the line on the screen.

I call this scenario *backwards causality*, because an action that can cause changes later in the signal transmission chain is taken to cause changes earlier in the signal transmission. Within the blend, the causal chain runs backwards.

4 Two Activities: Checking Impedances vs Quick Inserts

In this section, I discuss two of the diagrams that are used during the capping process. The diagrams are used to coordinate between artifacts with close to identical structure. In both cases, each electrode on the cap corresponds to a number on another artifact, the impedance meter or the connector box, arranged in a rectangular grid. The diagrams make these correspondences explicit. In contrast to their almost identical input structure, the diagrams look very different. I argue that the diagrams differ in the amount of integration between the input domains, and that the amount of interaction with the diagrams determines the extent of integration.

4.1 Checking Impedances

Recall that a major part of the capping process is the application of the electrode cap. All electrodes need to be prepared individually by lowering the impedance below the specified threshold. Pressing a numbered button on the impedance meter (Figure 2) displays the impedance for the corresponding electrode. A reading above threshold prompts the experimenter into a cycle of lowering the impedance at the electrode and checking the impedance at the impedance meter, with the purpose of bringing the impedance below threshold. Obviously, pressing the correct button is crucial, otherwise the experimenter's attempts at lowering the impedance will show no effect on the impedance meter. The numbered layout (Figure 3) displays the numbers of the impedance meter buttons at the locations of the electrodes, and serves as a mediating representation for the coordination between the cap and the impedance meter. It is frequently consulted during the capping process, both to look up the number corresponding to the electrode that was just worked on, and to make sure that the impedance meter number corresponds to the electrode that is being tested. Thus, an integral part of the capping process is the frequent coordination between cap electrodes and the impedance meter, and the numbered layout provides an integrated representation that allows the coordination of the electrodes on the cap and the buttons on the impedance meter.

The conceptual integration network with the input domains Electrode Cap and Impedance Meter is presented in Figure 11. The Electrode Cap provides the electrodes in a specific spatial arrangement. The Impedance Meter features numbered buttons in a different spatial arrangement. Electrodes and numbers are connected by a cross-space mapping. The two different spatial arrangements are not connected by a cross-space mapping, as they don't correspond to each other in any meaningful way. In contrast to the numbers and electrodes, which are both projected into the blend, only the spatial arrangement of the Electrode Cap is projected. The Electrode Cap and Impedance Meter domains are integrated by superimposing the numbers on the electrode positions in a diagram that preserves the topology of the arrangement of the electrodes in the cap.

The numbered layout serves as a material anchor for the blend of the Electrode Cap and the Impedance Meter. It is an external representation that gives stability to correspondences that can not be kept in memory without extensive training. The input domains are integrated in this representation: the button numbers from the Impedance Meter input are written into the circles representing the electrodes from the Electrode Cap, and are arranged within the layout of the cap.

4.2 Quick Inserts

The cap electrodes need to be connected to the amplifiers. This is normally accomplished by connecting the flat band cables of the cap with their corresponding counterparts in the connector boxes. However, when a quick insert was used to override a non-functional

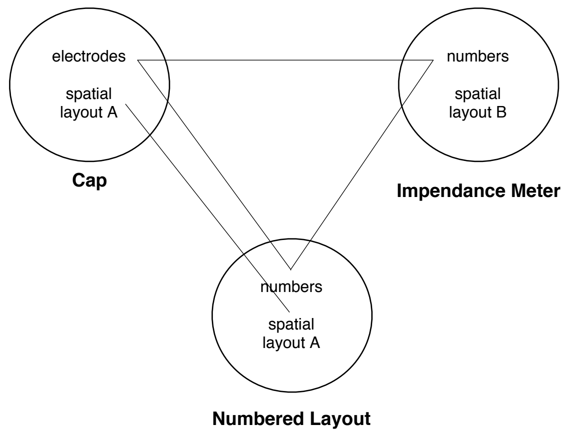

Figure 11: Checking Impedances

electrode, the experimenter needs to unplug the default connection from the flat band cable and replace it with the quick insert. The first two columns of the chart labeled “connectors for cap/ quick inserts” (Figure 6) gives the necessary information: the first column, labeled “CAP”, lists the electrode names, the second column, labeled “NUMBER”, refers to the corresponding number in the connector box (Figure 5). In order to replace the wire, the experimenter needs to look up the name of the electrode (using the labeled layout, Figure 4), find the label and the corresponding number in the chart, and then replace the electrode’s wire with the quick insert. The use of quick inserts during capping is an exception, so this task is only rarely performed.

The chart “connectors for cap/ quick inserts” combines information from two separate domains, the Electrode Cap and the Connector Box. The Electrode Cap is represented by the the electrode labels given in the labeled layout. The cap and the labeled layout both feature (labeled) electrodes in the same spatial layout. The connector box features numbers which are arranged in a square grid pattern on the connector box, and correspond to individual wires. The electrodes and the numbers on the connector box are connected via a cross-space mapping: the numbers stand for the wires which are connected to electrodes on the cap.

Unlike the numbered layout, the chart does not seem to reflect an integration of the two input domains. Neither of the input spaces provide their spatial layout for the chart. Instead, the correspondences of electrode labels and connector numbers are arranged in order of increasing number of the connection (column 2). Ordering a list in increasing order is a common ordering strategy, but this order is not particular to either of the inputs. Further, the information from the two domains is listed side by side in separate columns, not superimposed over each other. The chart thus represents the mapping between the labels and the numbers, not a blend between the domains. Figure 12 gives a schematic representation of the coordination of the input domains. The electrode labels correspond to the numbers in the connector box, and are thus linked by a cross space mapping. Labels and numbers are both part of the “connectors for cap/ quick inserts” chart, and thus projected, but neither of the spatial layouts is projected. The chart is represented as a rectangle instead of a circle, as a visual reminder that the two domains are not integrated.

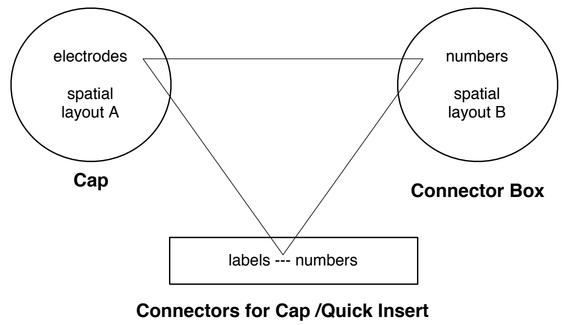

Figure 12: Quick Inserts

5 Discussion: The Relationship of Blending and Action

I presented several examples suggesting that conceptual blends play a role in everyday activity. Blends can trigger actions that are not meaningful in the input spaces alone. Backwards causality in the Brainwaves blend is the prime example: Actions later in a causal chain are taken to have an effect on events earlier in the causal chain. Since brain activity and its representation on the computer screen are compressed into uniqueness in the Brainwaves blend, actions that change the representation are taken to have an effect on the represented, the brain activity recording.

Backwards causality in the Brainwaves blend arises from the interplay of two principles of conceptual blending: topology and compression (Fauconnier & Turner, 2002). Preservation of topology of the input space mappings occurs when the cause-and-effect link between the reset button of the amplifier and the line on the screen is preserved in the blend: Pressing the button causes the line to go flat. Compression of cross-space mappings occurs when the electrode recording and the line on the screen are compressed into uniqueness in the blend: The line is the recording of brain activity. In conjunction, compression and preservation of topology introduce the novel cause-and-effect link between the electrode signal and the reset button into the blend: Only in the blend does the pressing of the reset button cause a change in the recording at the electrode. This novel cause-and-effect link is inconsequential in the input spaces. An interesting question is how expertise and training can change the interplay of the topology and compression principles in the conceptual government of activity.

Conceptual integration can guide activity, but, similarly, activity can determine the extent of the integration between two domains. Checking Impedances and plugging in Quick Inserts require the same sort of coordination: The electrodes on the cap need to be cross-referenced with an arbitrary set of numbers. The lab provides conversion charts for both. In the Checking Impedances example, the conversion chart displays the two domains in highly integrated form; this chart constitutes a material anchor for a blend between the two input domains. In the Quick Inserts example, the conversion chart shows no evidence of integration, and does not serve as a material anchor for a blend between the domains. Why are the domains integrated in one case but not the other?

The two examples strongly differ in the importance they have during the activity. The impedance meter is used constantly for at least 20 minutes. The numbered layout needs to allow cross-reference in both directions: from the cap to the impedance meter, and from the meter back to the cap. Presenting the numbers in the same spatial arrangement as the electrodes allows the identification of impedance meter button numbers in just a glance. The frequent use of the numbered layout makes its integrated form necessary for efficient capping.

In contrast, the connector boxes are used for only a short period of time, and they do not provide an output that is used in further action. The interaction with the connector box is usually unidirectional from the cap to the box: quick insert wires need to be plugged into the box. The opposite is conceivable: for instance, to double-check if the quick inserts are placed in the correct position, one might read the number on the connector box with the quick insert wire, look up the corresponding electrode, and check whether the quick insert is located in that electrode. It is an empirical question whether and with which frequency this occurs. In the normal case of events the interaction with the connector box and the "connectors for cap/ quick inserts" is so restricted that an integration of the two domains, Connectors and Cap, would not make the overall process any more efficient. In this situation, an integration is not needed.

Differences in activity and in the extent of interaction with artifacts can thus lead to completely different combination of identical input structures. Extensive interaction leads to strong integration, minimal interaction to minimal integration. Blends are thus tied in tightly with the activity in the lab.

In conclusion, blends can give rise to actions that are not meaningful when considered in separate input spaces alone. Blending results from action: a blend is likely to be constructed if it makes activities more efficient. Conversely, if the activity is not important enough within the overall task, integration is not needed. The presence and absence of integration is reflected in the mediating representations that make the correspondences between the artifacts in the lab explicit.

References

Coulson, S., & Oakley, T. (2000). Blending Basics. *Cognitive Linguistics*, *11*(3/4), 175-196.

Fauconnier, G., & Turner, M. (2002). *The Way We Think.* New York: Basic Books.

Hutchins, E. (2002). Material Anchors for Conceptual Blends. In *Proceedings of The Way We Think Symposium on Conceptual Blending.* University of Southern Denmark, Odense.

Cognitive Science Online, Vol.1, pp.46–57, 2003 http://cogsci-online.ucsd.edu

Lexical dynamics and conceptual change: Analyses and implications for information retrieval

Robert Liebscher & Richard K. Belew
Department of Cognitive Science
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0515
rliebsch|rik@cogsci.ucsd.edu

Abstract

One important aspect of a document's context is the time at which it was written. We report here on analyses of formal (dissertation abstracts) and informal (discussion board postings) communications among academics within two separate disciplines. We focus on academic communications because these especially must be understood within the context of what has been said before, together with what is considered relevant and worth saying at the time of publication. All corpora include time-stamp information that allows temporal analysis of changing lexical frequencies across decades. Using techniques borrowed from time series analysis, we find distinct patterns of "rising" and "falling" bigram frequencies in both domains, and argue that this information can be exploited to improve retrieval of relevant documents.

1 Introduction

A common idealization of many information retrieval (IR) tasks is to reduce retrievable documents and queries to simply the sets of lexical keywords they contain. But it is becoming increasingly acknowledged that if IR systems are ever to qualitatively improve their ability to help users seek relevant documents, increased attention must be paid to the *context* shaping use of these keywords by both the browsing user and the documents' original authors. In many situations, one obvious feature of both users' and authors' contexts is their respective places in time. As typically addressed, the retrieval task makes almost no assumptions about time. In general, the time frame in which an author commits her thoughts to writing, as well as the time frame within which the querying user operates, are either assumed to be irrelevant to effective retrieval, or tacitly assumed to be roughly contemporaneous with one another.

But the conceptual frameworks within which authors write can play an enormous role in what they choose to say, and how they choose to say it. This is especially true in scientific writings, where acceptance by peer reviewers is essential to success and recognition. By the same token, scientists and students searching historically through the body of work of a science will often be familiar with a different vocabulary as they pursue a more modern agenda.

Within the field of IR system design, our reflection of the changing semantic understanding of a science is potentially informed by observation of word frequency statistics combined with metadata capturing the document's publication time. Words and phrases change in both frequency of use and meaning through time. For example, the token `WEB` is found in many areas of discourse today, though little more than a decade ago, prior to the advent of the World Wide Web, its use was much more circumscribed.

In this paper, it is assumed that within a given domain, the overall frequency of a term k at time t is proportional to a community's collective "interest" in k at t. As interest in a topic waxes or wanes, the raw *number* of documents containing information about that topic rises or falls through time. These arguments, along with observations about the nature of scientific discourse, are used to construct a temporal weighting scheme that places a document in an appropriate historical context.

Consider a "rising" term k, which moves from obscurity to immense popularity over the duration of a corpus. In an academic context, it is not unreasonable to assume that the term was once part of a small group of "seminal papers" that helped to launch a field of inquiry, a technology, a methodology, etc. (To be concrete, imagine a search for the now ubiquitous term `DNA`. We would surely want Watson & Crick's one-page paper of 1953 to be deemed relevant!) Under an atemporal paradigm, a query for k will return a temporally random subset of documents in the corpus, leaving our user in the dark with regard to any notion of the conceptual development of her query term. The seminal papers have a chance of being deemed relevant that is proportional only to their length.

But with the temporal weighting scheme introduced in Section 3, when the frequency of this rising term k is initially low (i.e. used in very few documents), its weight will be amplified. At a later point in time, when its use is much more common, its weight will be dampened so as not to over-emphasize the many documents about k that exist at that time.

Alternatively, imagine a "dying" term, where k is omnipresent at one point in time, then falls in frequency until it is no longer used. Under a temporal weighting scheme, its initial use is dampened, and its later use is amplified. One consequence of this would be to emphasize *historical* documents that are written retrospectively about the term in question. These provide a good starting place for someone who wishes, as above, to gain an historical perspective on the development of k.

This paper reports attempts to formalize these arguments and improve information retrieval by incorporating the time at which a document was written into the retrieval process.

2 Methods

2.1 Corpora

The analysis will concentrate on abstracts from Doctoral and Masters' dissertations because these are available across decades in a relatively consistent format. Further, the focus is placed on a particular discipline, artificial intelligence, in order to relate the analysis of changing frequency statistics to the semantics of the evolving science that generates them.

While this is terrifically rich data, the amount of text provided by only dissertations' titles and abstracts is not great. This is especially unfortunate, because standard time series analysis demands large data volumes. However, as will be demonstrated, much information can be gained from an analysis that employs the most straightforward linear models of trend.

Three corpora were used in these studies. The first, *AIT*, contains approximately 5,000 abstracts from Ph.D. and Masters theses in artificial intelligence, collected by University Microfilms, Inc. from 1986 to 1997 (Belew, 2000). Each document is labeled with its year of publication.

The second corpus, *CommDis*, also from University Microfilms, Inc. contains approximately 4,000 abstracts from Ph.D. and Masters theses in language and communicative disorders. The abstracts run from 1980 to 2002, and each is labeled with its year of publication.

The third corpus, *AIList Digest* (hereafter AIList), is a subscription-based electronic newsletter that contains over 10,000 discussion board postings, conference announcements, and essays on artificial intelligence that were collected and distributed weekly from 1983 to 1988. Each document is labeled with the exact time and date on which it was written. As the intent of this study was to treat AIList as a record of informal academic communication, some documents, such as bibliographies and subscription statistics, needed to be removed. A very simple type/token ratio filter worked well to preserve relevant articles of discourse while filtering out the unwanted documents, which generally contained few grammatical terms relative to the total number of tokens.

2.2 Tracking lexical frequency change

Bigram frequency counts were made for each of the 12 years in AIT, for each of 12 bins of roughly equal size (approximately 110,000 tokens each) in AIList, and for each of the 23 years of CommDis. Statistics were also collected on unigrams, but for the purposes of this work, bigrams provide a much greater level of detail and are considered better descriptors (and therefore more likely to serve as query terms) in domain-specific corpora (Damerau, 1993). Unigrams of particular interest are abbreviations and acronyms, which are discussed in Section 5.

The task of modeling the frequencies of terms that have consistently increased or decreased through time lends itself well to formal methods in time series analysis (Box et al., 1994). However, the 12 data points associated with AIT and AIList are too sparse to allow accurate modeling, and even the 23 points associated with CommDis only approach data sufficiency. For this reason, the analysis here is restricted to simple *linear* models of trend.

Mean smoothing was performed over the temporal frequency plot of each term that met a minimum frequency requirement, each bin being averaged with its two neighbors. Smoothing is especially necessary for AIList, as an extended thread of conversation or long essay might cause a spike in the frequency of a particular term that does not accurately reflect the level of community interest at that time.

Each smoothed frequency plot was then fit with a regression line, and the adjusted correlation coefficient (between time and frequency) r and slope of the line b were measured. The slope b of a term's temporal frequency plot is only a meaningful number when compared to the slopes of other terms. Terms which steadily increased in frequency through time will have positive slopes; those which steadily decreased will have negative slopes. Terms with a slope of $2s$ rose (or fell) twice as quickly as those with a slope of s.

Figure 1 shows the temporal frequency plot of two examples drawn from each corpus. Terms that met a minimum threshold for r (0.70) and absolute value of b (corpus-specific)[1] were extracted for further study. These included 49%, 42%, and 28% of the bigrams, respectively, for AIT, AIList, and CommDis, and represent the terms that might benefit from a temporal analysis of frequency.

Table 1 depicts for the top ten "rising" and "falling" terms extracted from the AIT corpus, along with their b and r values. While the frequencies of many terms within a domain

[1]Careful manual examination of the data led to the choices of 5.0, 4.0, and 7.0 for threshold values of b for AIT, AIList, and CommDis, respectively.

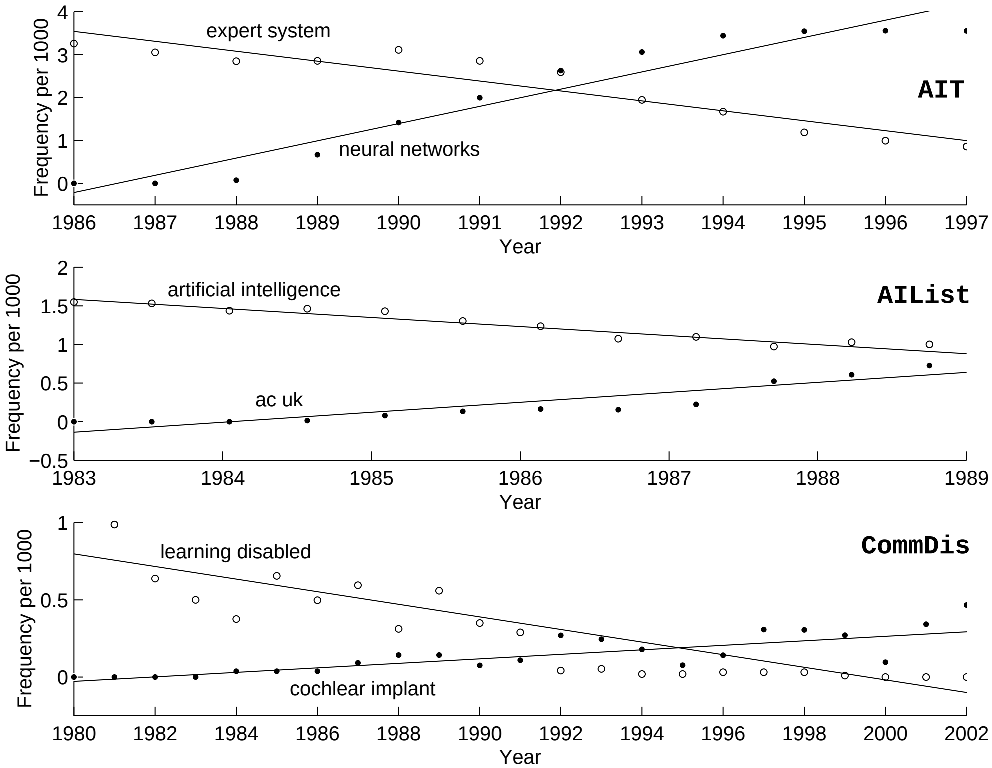

Figure 1: Examples of rising and falling terms from the three different corpora, with regression lines. Top to bottom: AIT, AIList, CommDis.

are temporally invariant or random, there are many others that undergo large frequency changes over time. In the traditional retrieval task, this information is not exploited.

Appendix A contains similar tables for AIList and CommDis. The results from CommDis dissertations provide strong confirmation that the analysis of conceptual change within the domain of AI transfers well to this new domain, despite large differences in the time frame and rate of dissertation publication across the two corpora. The results of the AIList newsgroup analysis are more mixed. While some of AIList's changing bigrams do indeed overlap semantically with those found in AIT, others appear to be "noise." These are most likely a consequence of both small AIList corpus size and difficulties in cleanly parsing the highly variable news postings (e.g., identifying common "signature lines" used by frequent posters). These anomolies may, however, also point to qualitative features of informal language use within such groups that limit the utility of the methods being used.

3 Temporal term weighting

The simplest form of **TF-IDF** weighting multiplies the raw **term frequency** (TF) of a term in a document by the term's **inverse document frequency** (IDF) weight:

$$idf_k = \log\left(\frac{NDoc}{D_k}\right) \tag{1}$$

Term	Slope (*b*)	*r* value
NEURAL NETWORK	474.32	0.9283
NEURAL NETWORKS	384.24	0.9505
FUZZY LOGIC	120.88	0.9035
ARTIFICIAL NEURAL	95.15	0.8990
GENETIC ALGORITHMS	55.14	0.9624
REAL WORLD	42.31	0.8509
REINFORCEMENT LEARNING	36.20	0.8447
PATTERN RECOGNITION	35.71	0.9314
NONLINEAR SYSTEMS	32.50	0.9511
LEARNING ALGORITHM	32.17	0.8313
ARTIFICIAL INTELLIGENCE	-248.07	-0.9309
EXPERT SYSTEM	-222.33	-0.9241
EXPERT SYSTEMS	-194.42	-0.9769
KNOWLEDGE BASED	-125.99	-0.9144
PROBLEM SOLVING	-90.48	-0.9490
KNOWLEDGE BASE	-73.65	-0.9281
KNOWLEDGE REPRESENTATION	-61.18	-0.9603
DECISION MAKING	-51.31	-0.9521
BASED EXPERT	-43.08	-0.9846
RULE BASED	-41.74	-0.9636

Table 1: Top ten rising and falling bigrams from AIT (1986-1997). Informal queries of AI practitioners revealed that the terms in these lists matched well with their memory of developments in the field over the period in question.

$$w_{kd} = f_{kd} \cdot idf_k \tag{2}$$

where f_{kd} is the frequency with which keyword k occurs in document d, $NDoc$ is the total number of documents in the corpus, and D_k is the number of documents containing keyword k.

In a temporal model, the corpus is divided into a series of independent sub-corpora, each associated with documents occurring within a particular time slice. IDF weights can then be computed independently for each sub-corpus.

One way to characterize the change is to contrast the temporally changing weights we propose as a *difference* with the traditional IDF weighting:

$$\begin{aligned} \Delta idf &= \log\left(\frac{NDoc_k(t)}{D_k(t)}\right) - \log\left(\frac{NDoc}{D_k}\right) && (3) \\ &= \log\left(\frac{\frac{NDoc_k(t)}{D_k(t)}}{\frac{NDoc}{D_k}}\right) && (4) \\ &\simeq \log\left(D_k/D_k(t)\right) && (5) \end{aligned}$$

where $D_k(t)$ is the number of documents in time slice t that contain k. Figure 2 shows how a keyword that changes in frequency over time can effect weighting. Following the assumption of a linear model in Section 2.2, a linearly increasing keyword is also assumed here. Note that the changes to the multiplicative IDF factor swing dramatically. With a keyword that increases in frequency over time, early uses of the term cause positive contributions, while later uses produce decreased contributions. The same argument holds for a term that decreases in frequency over time.

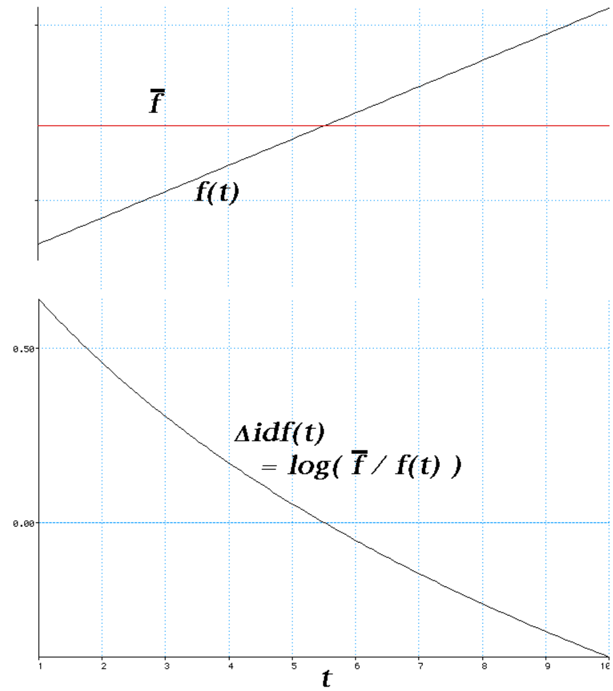

Figure 2: Dynamic IDF

Corpus	$\overline{r}(TF(t), D(t))$	$\overline{r}(TF(t), f(t))$
AIT	0.8899	0.3038
AIList	0.8885	0.3383
CommDis	0.8175	0.4007

Table 2: Correlation coefficent ($\overline{r}$) values for the three corpora.

Two assumptions must be made explicit before proceeding. First, f_{kd}, the number of times a term k occurs in document d, is only expected to correlate with the document's length. The average number of times that a document of length l mentions a term k, then, does not vary with respect to time. Second, documents of the same type within a domain do not become longer or shorter, on average, over time.

If these assumptions are true, then $TF_k(t)$, the total frequency of k in time slice t, should be proportional to $D_k(t)$, a rough measure of collective community interest in k at time t. This means that the following should hold:

$$log(D_k/D_k(t)) \simeq log(TF_k/TF_k(t)) \tag{6}$$

Table 2 shows that the number of documents in which a term appears at time t can be approximated by the total term frequency, as measured by the correlation coefficient $\overline{r}$ averaged over all terms. Note the striking similarity between $TF_k(t)$ and $D_k(t)$. Furthermore, the correlation between total term frequency and within-document frequency is low[2], supporting our first assumption above. f_{kd} is a constant over all t, as the burden of "temporal

[2] $TF(t)$ and $f(t)$ are not entirely *un*correlated ($\overline{r} = 0.0$) because there are some very infrequent terms in each corpus that have one document written specifically about them. These terms spike in

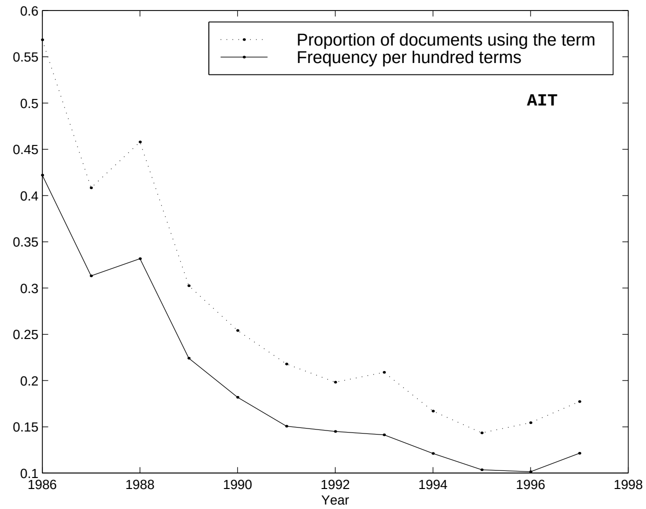

Figure 3: $TF_k(t)$ and $D_k(t)$ for the term `ARTIFICIAL INTELLIGENCE` over the duration of the AIT corpus. $TF_k(t)$ is frequency per hundred terms for scaling purposes.

culpability" in determining $TF_k(t)$ falls squarely on the weighting factor it approximates, $D_k(t)$. Figure 3 demonstrates this for the term `ARTIFICIAL INTELLIGENCE`; this example is discussed in greater depth in Section 5.

One virtue of using the techniques of time series analysis is that an estimate of what temporal perturbation is appropriate with respect to a particular document benefits from the more stable statistics of the entire series. Further, the fact that simple parametric models are fit means that these new weighting factors can be computed sufficiently at run/retrieval time. Linear models may be most appropriate for the corpora used in this study, but the same arguments apply to more elaborate (e.g. exponential; see Section 4.2) models of lexical frequency change as well.

4 Related work

4.1 Topic Detection and Tracking

Within the IR community, Topic Detection and Tracking (TDT) represents the most focused attention to special features in dynamic textual corpora. Research in TDT has focused primarily on the identification of *events* as reported in the news. Authors (e.g., newswire reporters) generate text capturing their reactions to, and interpretations of, "external" events arising in the world. There are obviously many commercial and security applications which motivate the specific TDT applications (e.g. story segmentation, new event detection, and event monitoring).

In part, the present work is motivated by the recognition within the TDT community that these methods must be extended:

> The results that we have presented [in this review paper] on the three

within-document frequency in the bin containing their signature document, while remaining flat over the rest of the corpus, and this is reflected in r. This is not very useful here, as the aim is to find terms that follow nonrandom temporal trends.

> detection tasks were acceptable, but not as high a quality that we would have liked. We believe that we have hit the limits of the effectiveness that can be reached with simple IR based approaches to story/topic comparison (Allan et al., 2002).

The approach of the current study has been to attempt to model patterns of change in the lexicon across time, and then more directly introduce the consequences of these changes. To this end, it is reasonable to believe that the methodology of TDT can be extended by considering the sociocultural contexts in which documents are written.

There are several ways in which the data used in TDT can be expected to differ from that in the corpora introduced in this study. Most obviously, the time scale differs dramatically. The AIT, AIList, and CommDis corpora contain documents written over the course of decades, while (for example) the TDT-2 corpus is collected over the first six months of 1998. Also, the TDT-2 corpus benefits from relevance assessment labels from 200 different topics on individual stories, whereas no such labels exist for the individual documents in the artificial intelligence and communication disorders corpora.

But there are more subtle factors influencing changes in the corpora as well. The *conceptual movement* of central interest in scientific discourse can be argued to be much more internally driven by the literature itself. Certainly there are "events" that ripple through scientific disciplines, but the amount of analysis, interpretation and theorizing is considerably greater. Scientific discoveries, when transformed into text, help to establish the next line of inquiry (Latour, 1986). In this way, the real world and the corpora continually influence one another. This places limits on a purely statistical approach to studying the corpora.

4.2 Recent models of trend

An interesting attempt within the TDT community to examine this dialectic is a study of the relations between financial news stories and stock prices by Lavrenko et al. (2000). They used piecewise linear regression to identify trends in a stock's price, then constructed language models from stories about the stock that preceeded gains or declines. The models were used to train a system to predict changes in price based upon the words used in an incoming stream of news stories. As investors spend a significant amount of time researching potential trades, as well as watching the fluctuations in a stock's price, the mutual influence of text and the outside world is obvious. Their system significantly outperformed a random trading simulation.

A key difference between Lavrenko et al. (2000) and the present work is, once again, time scale. As Web-based financial publishing and stock trading happen on the scale of minutes and hours, compared to months and years for scientific publishing and experimentation, there may not be reason to expect similar methods to apply in discovering trends. Piecewise linear segmentation is necessary for the financial study, and will likely play a role in the future work discussed in Section 5. In the present work, many terms followed a fairly linear increasing or decreasing trend in frequency, and these trends are likely a by-product of the particular time periods covered by the corpora. Many terms rose and fell over the duration of a corpus (particularly CommDis, which spans 23 years), and as such were accorded a poor goodness-of-fit to a single line spanning the corpus' length.

A recent attempt to formulate a methodology that works on any time scale is that of Kleinberg (2002). His model posits an infinite state automaton that takes advantage of the observation that, independent of context or type of discourse, terms generally occur in "bursts". That is, when a term is seen, its likelihood to appear again soon increases. Informally, the model can be described as performing piecewise exponential curve fitting. In one of several examples tested with the model, terms from conference paper titles spanning several decades were identified during periods of exponential growth in frequency as a way to find the terms "... that exhibit the most prominent rising and falling pattern over a limited period of time."

Other cases examined in Kleinberg (2002) include modeling arrival times of e-mails containing a particular term concerned with an event in the world (such as the deadline for a grant proposal) and modeling term usage in U.S. Presidential State of the Union addresses over the course of two centuries. Besides spanning different time frames, another advantage of Kleinberg's model is its independence from requirements of formal time series analysis. It was mentioned earlier that the corpora used in the present study reveal a data sufficiency problem, when viewed as a problem for standard ARIMA-style modeling (Box et al., 1994).

Given this, a comparison of piecewise linear and exponential approaches to discovering trends in term usage seems in order. The main argument against a linear approach is data sufficiency, while a potential argument against an exponential model is that, while it is able to handle bursts of increasing intensity, it is difficult to imagine (outside of news corpora) the possibility of "reverse burstiness", i.e. a term exhibiting exponentially decreasing frequency. These are empirical matters to be investigated, along with the others described in the following section.

5 Conclusions and future work

Benefits to information retrieval, particularly for a user who is inexperienced in a particular domain, can result from paying attention to changes in term frequencies through time. In both formal and informal scientific communications, simple linear trends fit the temporal frequency curves for many terms, some of which showed dramatic rises or falls. Based upon the assertion that an academic community's "collective interest" in a topic at a given time is proportional to the frequency at which the topic is mentioned, a temporally-sensitive alteration to the standard (atemporal) TF-IDF weighting scheme was suggested. This allows documents, particularly within academic communication, to be placed in their proper historical context. The recent availability of collections of conference proceedings such as ACL and ACM-SIGIR promise to provide much more complete evidence, and allow more elaborate models.

Kleinberg has commented with regard to the burstiness of terms, and indeed most events, that our memory is structured with respect to these bursts in such a way that "... particular events are signaled by a compression of the time-sense" (Kleinberg, 2002). One interpretation of this statement is that readers recall historical periods in terms of a *Zeitgeist*, not a fluid timeline. For example, 1982 is sometimes referred to as the "year of the computer", when it was named Time Magazine's "Person of the Year"; similarly, the 1990s are sometimes described retrospectively as the "decade of the brain". Short of capturing features of popular culture's *Zeitgeist*, a potentially interesting way to create a *gold standard* for studies such as these is to collect data from scientific participants in the fields of artificial intelligence, communicative disorders, and (after the ACL corpus has been analyzed) computational linguistics. Future experiments will investigate the extent to which the present analyses are consistent with their historical recollections of the fields.

The behavior of the term `ARTIFICIAL INTELLIGENCE` in the AIT and AIList corpora suggests other linguistic and sociocultural interpretations of the data. One contributing factor seems to be the "substitution" phenomenon of the full phrase `ARTIFICIAL INTELLIGENCE` with its *acronym* `AI`. Within the AIList corpus, the decrease in frequency of `ARTIFICIAL INTELLIGENCE` is accompanied by a rise of the acronym `AI`, which was the sixth fastest rising unigram. Similarly, `COMPUTER SCIENCE` experienced the sixth fastest decrease, while `CS` was the thirteenth fastest riser.

The same patterns do not hold true in AIT, perhaps suggesting an important difference between informal (AIList discussion board postings) and formal (AIT dissertation abstracts) as channels of academic communication. Schwartz and Hearst (2003) have recently developed a fast algorithm for identifying abbreviations and acronyms and their full-phrase referents in the biomedical literature, which is presently experiencing explosive growth both in volume and coinages of new terms and abbreviations. A temporal examination of the rise of different acronyms and fall of their referents in different domains is an obvious next step.

Another contributing factor may involve the phenomenon of "internal vs. external keywords." Speaking of a token frequency analysis of the same AIT corpus:

> ... the statistics for stemmed, non-noise word tokens, noise words (e.g. `the`) are [combined]. As expected, the noise words are very frequent. But it is interesting to contrast those very frequent words defined *a priori* in the negative dictionary with those that are especially frequent in this particular [AIT] corpus. In many ways these are excellent candidates for *external keywords*: characteristizations of this corpus' content, from the "external" perspective of general language use. That is, these are exactly the words (cf. `NEURAL NETWORK, BASE, LEARN, WORLD, KNOWLEDGE`) that could suggest to a browsing WWW user that the AIT corpus might be worth visiting. Once "inside" the topical domain of AI, however, these same words become as ineffective as other noise-words, as *internal keywords*, discriminating the contents of one AI thesis dissertation from the next (cf. `SYSTEM, MODEL, PROCESS, DESIGN`). (Belew, 2000, pp. 71,72)

Considered from a temporal perspective, it seems that after a period of using the term, the practitioners of artificial intelligence began to assume its use to be *tacit* at least in some cases, and as such did not make explicit reference to it. Just how long this process takes, which factors influence this rate, etc., all become potentially interesting questions.

Also at work is the phenomenon of lexical replacement not by an acronym but by another term that, in some situations, becomes interchangeable with the original term (Bauer, 1994; Howard, 1977). As `ARTIFICIAL INTELLIGENCE` fell in AIList and AIT, `MACHINE LEARNING` rose in both corpora, and in fact was the eleventh fastest rising term in AIT. A more conspicuous example comes from the CommDis corpus, where the fastest falling unigram was `SUBJECT`, while the fastest riser was `PARTICIPANT`.

One can also imagine this analysis being extended to consider the differential flow of conceptual markers through various publication forums. For example, many of the same rising and falling bigrams identified within the more formal and chronologically later AIT dissertation corpus seem to be *anticipated* by similar usage patterns within the more informal and earlier AIList newsgroups. While it seems reasonable to expect such linguistic patterns to be "worked out" within informal contexts prior to more formal adoption by a community, insufficient data with respect to the AIT and AIList corpora make it premature to reach this conclusion.

Looking ahead still further, investigations of *groups* of terms that rise and fall together in time, as well as conceptual movements among such systems of vocabulary, are in order. Combined with holistic analyses of entire vocabularies (e.g., eigen-structure analyses such as Latent Semantic Indexing (Deerwester et al., 1990)), such methods may provide some of the most solid empirical foundations for an understanding of conceptual change within scientific communities. Almost certainly, knowledge about the nature of language change should support improved performance by temporally-informed retrieval and classification systems.

Term	Slope (*b*)	*r* value
AILIST DIGEST	206.65	0.9535
VOLUME ISSUE	142.95	0.9762
AC UK	92.82	0.9371
MIT EDU	81.99	0.9470
NEURAL NETWORKS	52.33	0.8482
STANFORD EDU	42.35	0.9798
NEURAL NETWORK	41.54	0.8879
NEURAL NETS	33.52	0.8580
AI AI	32.50	0.8302
AI MIT	32.07	0.8971
NATURAL LANGUAGE	-85.41	-0.9402
EXPERT SYSTEMS	-71.76	-0.6261
ARTIFICIAL INTELLIGENCE	-71.34	-0.8766
EXPERT SYSTEM	-58.61	-0.3878
LOGIC PROGRAMMING	-54.87	-0.5980
COMPUTER SCIENCE	-52.78	-0.8242
SRI AI	-51.80	-0.9624
TURBO PROLOG	-37.48	-0.8752
PROGRAMMING LANGUAGE	-36.12	-0.7959
CSNET RELAY	-35.88	-0.9629

Table 3: Top 10 rising and falling terms in the AIList corpus.

Term	Slope (*b*)	*r* value
HEARING LOSS	60.5875	0.849
PHONOLOGICAL AWARENESS	43.2302	0.865
WORKING MEMORY	38.9826	0.951
HEARING AID	38.8868	0.456
SPEECH LANGUAGE	25.4638	0.490
CHILDREN SLI	23.1939	0.573
FORMANT TRANSITIONS	20.4096	0.613
HEALTH CARE	19.7127	0.611
SPEECH RECOGNITION	18.5384	0.651
COCHLEAR IMPLANTS	18.1522	0.831
HEARING IMPAIRED	-56.6631	-0.837
SB RM	-46.9355	-0.711
LANGUAGE IMPAIRED	-43.5814	-0.896
LEARNING DISABLED	-43.2694	-0.882
NORMAL HEARING	-40.7301	-0.881
CLOSURE DURATION	-34.1147	-0.940
IMPAIRED CHILDREN	-31.0231	-0.856
FUNDAMENTAL FREQUENCY	-30.6432	-0.771
ACOUSTIC REFLEX	-30.0184	-0.797
DISABLED CHILDREN	-29.3471	-0.716

Table 4: Top 10 rising and falling terms in the CommDis corpus. Note that some of the falling terms are as much a product of “political correctness” as they are of the nature of inquiry in the field.

References

Allan, J., Lavrenko, V., & Swan, R. (2002). Explorations within topic tracking and detection. In J. Allan (Ed.), *Topic detection and tracking: Event-based information organization* (p. 197-224). Kluwer Press.

Bauer, L. (1994). *Watching english change.* London: Longman Press.

Belew, R. K. (2000). *Finding out about: A cognitive perspective on search engine technology and the world wide web.* Cambridge University Press.

Box, G., Jenkins, G., & Reinsel, G. (1994). *Time series analysis: Forecasting and control.* Prentice Hall.

Damerau, F. (1993). Generating and evaluating domain-oriented multi-word terms from texts. *Information Processing and Management*, *29*(4), 433-447.

Deerwester, S., Dumais, S., Furnas, G., Landauer, T., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society of Information Science*, *41*(6), 391-407.

Howard, P. (1977). *New words for old.* Hamish Hamilton.

Kleinberg, J. (2002). Bursty and hierarchical structure in streams. In *8th ACM SIGKDD international conference on knowledge discovery and data mining.*

Latour, B. (1986). Visualization and cognition: Thinking with eyes and hands. In *Knowledge and society: Studies in the sociology of culture past and present* (Vol. 6, p. 1-40). JAI Press, Inc. (transformation of rats and chemicals into paper.)

Lavrenko, V., Schmill, M., Lawrie, D., Ogilvie, P., Jensen, D., & Allan, J. (2000). Mining of concurrent text and time series. In *6th ACM SIGKDD international conference on knowledge discovery and data mining, text mining workshop* (p. 37-44).

Schwartz, A., & Hearst, M. (2003). A simple algorithm for identifying abbreviation definitions in biomedical text. In *Proceedings of the Pacific Symposium on Biocomputing.*