

Brainstorm: A High Level Overview

Jett Hays

Human Computer Interaction Department

Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA

jehays@andrew.cmu.edu

What

Here are a few metaphors to help you understand what Brainstorm *is*:

Brainstorm is a map that tells you how to get from neural signals to language.

Brainstorm is a set of instructions that tells you how to assemble language from thought.

Brainstorm is a brain to English translator.

Brainstorm abstracts what it does with a neat user interface that can run on any personal computer. The result of a Brainstorm session is a chunk of text (created with your brainwaves) that can be shared with the world through email. Brainstorm is also a development platform that allows researchers to rapidly test their own mind controlled keyboard ideas.

Why


Effective communication systems— from Twitter to the telephone— reduce the friction between what we think and what we share. Existing systems like keyboards do a good job of this, but they are still rudimentary. Even the most advanced interfaces like touchscreens are equivalent to punching a bunch of different buttons. By interacting directly with the brain, mind controlled keyboards can expand the intellectual throughput of our society. Brainstorm represents a significant step in the long journey to typing with our minds.

Definitions

For those unfamiliar with the technical terms, here are a few definitions:

EEG: A noninvasive tool to measure brain activity. A series of sensors measure changes in voltage along the brain's surface. The image below shows Brainstorm's sensor arrangement.

Brainstorm Sensor Positions



An overhead view of the Brainstorm's sensor positions. Each sensor produces a unique value.

Natural Language Processing: A branch of machine learning that makes predictions about text.

Frequency Band: A range of frequencies (how often something occurs), often used to describe complex time series data.

Features: Information that can be used to make a prediction.


Logistic Regression: An algorithm used by Brainstorm to classify EEG data. Each classification includes a 'confidence' score that indicates the likelihood of a given class.

Challenge

EEG Data is notoriously noisy. Determining what someone is thinking [the signal] requires filtering out irrelevant information [noise]. This process follows a long and intricate sequence of steps that uses complex algorithms. However, as you'll see, understanding *how* Brainstorm gets from no signal to noisy signal to a decision about

what you're thinking is easy once you consider the inputs, outputs, and purpose of each component.

Example Brainwave



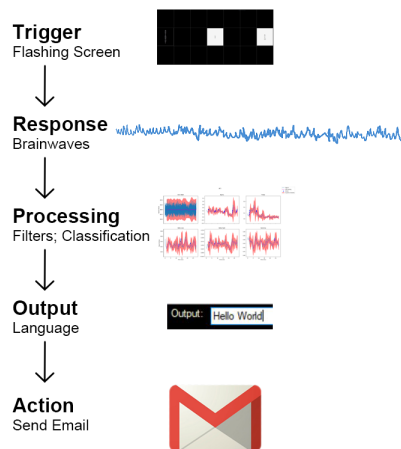
An example brainwave. As you can see, the eeg reading goes up and down a lot with no obvious pattern.

Solution

“Any fool can know. The point is to understand.” — Albert Einstein

Brainstorm uses a combination of techniques— including signal transformations, natural language processing, and machine learning— as it runs through the five step process shown below.

Brainstorm System Flow



The five main events Brainstorm initiates to translate brainwaves to text.

While you read descriptions of the following steps and their purpose, keep in mind the ultimate goal of Brainstorm: to bridge thought and language with a noninvasive brain computer interface. Every line of code and each facet of Brainstorm's design accomplishes a part of this goal. When combined into one system, these parts become a complicated whole. But if you can hold the purpose of Brainstorm in one hand, while you hold technical details in the other, you can grasp how Brainstorm translates brainwaves to language.


1. Trigger

Key Point: Brainstorm initiates the SSVEP phenomenon with boxes flashing at specific frequencies.

If you were to guess what someone was thinking at a random point in time, how would you do it? There are a number of different tactics you could adopt, but each approach relies on gathering additional information. The information gathering process could use a number of different methods— analyzing prior writing for thought patterns, reading body language, etc.— but every method is dependent on behavior. For decades, this was how scientists gathered information about the brain: by using behavior as a proxy for thought. But the development of brain imaging technology has allowed us to study what happens between thinking and doing. However, the brain is incredibly complex and the tools we have developed to study it are still rudimentary.

Unless you're willing to cut open the skull to insert a bundle of expensive sensors, there is a tradeoff between spatial and temporal resolution. This may sound complicated, but it really just means that the more certain we are of *where* something happened in the brain, the less certain we are of *when* something happened, and vice versa.

Time-Space Resolution Tradeoff



A simplified illustration of the tradeoff between our knowledge of where and when something occurs.

Brainstorm uses a wireless EEG headset. While noninvasive and relatively inexpensive, the EEG has poor *spatial* resolution. That means Brainstorm is limited to analyzing general regions of the brain, as opposed to specific neurons. In short, the data Brainstorm must process is frequent [refreshing 128 times per second], but imprecise.

To compensate for this uncertainty, Brainstorm controls a user's environment. As shown in the video below, Brainstorm flashes a series of 'language boxes' at specific frequencies. By constraining the problem to determining what frequency a person is looking at, Brainstorm is able to take advantage of the SSVEP phenomenon.

2. Response

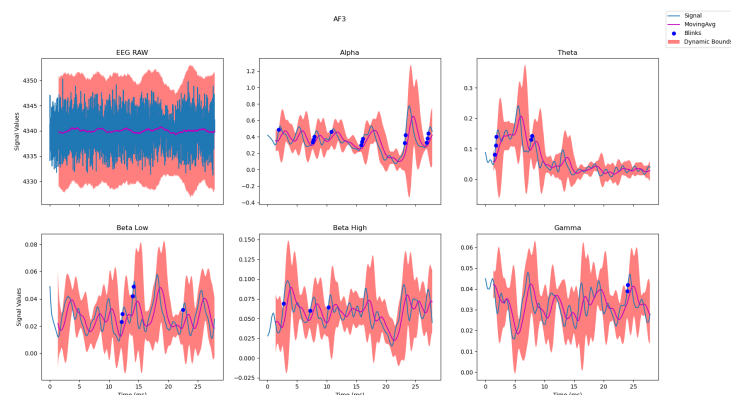
Key Point: The SSVEP phenomenon is when a flashing object induces neural activity at the same frequency. Brainstorm receives 8960 new pieces of information about this interaction between trigger and response every second through its 14 electrodes.

The SSVEP phenomenon is very simple: frequencies between 3.5 Hertz and 75 Hertz induce similar frequencies in the brain. Scientists like to use this phenomenon when studying neural activity, as SSVEP is consistent across people and resistant to artifacts.

An easy way to understand the SSVEP phenomenon is to think about a hall of mirrors. With the hall of mirrors analogy, the trigger [flashing boxes] is in the center of a room surrounded by mirrors. Every time the trigger fires, the mirrors bounce the signal back and forth through the room. As the signal bounces through the room, it becomes slightly distorted. Instead of a clear picture of the trigger, we are left with dozens of mirrors each containing a distinct portion of the original signal.

Instead of mirrors, Brainstorm uses 14 electrodes to 'reflect' the brain waves traveling through the room [a brain]. Each electrode produces five frequency bands that, when put together with machine learning, produce a clear picture of the trigger. This picture is assembled across time, with 8960 new data points produced every second. The filtered output from one electrode is shown below.

Sample Electrode Output



The data received from one electrode across a recording session.

3. Processing

Key Point: Brainstorm uses logistic regression for SSVEP classification. Feature selection, binning, and majority voting increase the accuracy of each classification.

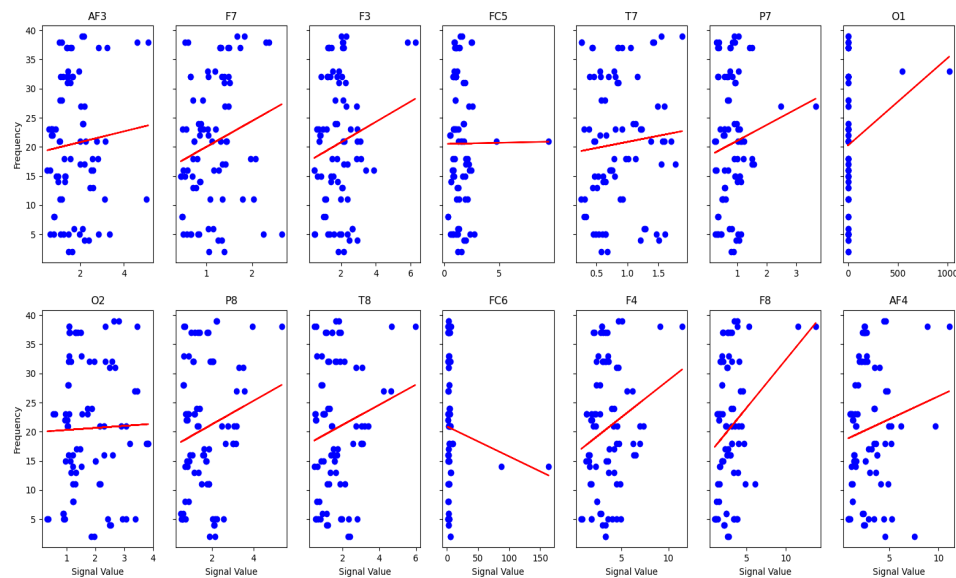
To begin, consider what Brainstorm's prediction system is being asked to do: *given 70 numbers, return the most likely frequency.* Satisfying this request, requires Brainstorm to balance long term patterns with short term noise. Brainstorm uses the following techniques to achieve this balance.

Band Selection

The brain is like a room, and neurons are like people. Groups of people chat and the room fills with noise. Joining a worthwhile conversation is hard, but if you visit enough groups you can gain a better understanding of which conversations are relevant and which are not.

With Brainstorm, groups of neurons are represented by electrodes. Each electrode has a different 'conversation', but only some of these conversations provide meaningful information about the SSVEP paradigm. The chart below shows the association between the stimulus and the strength of an electrode's response, which provides an indication of which electrodes provide meaningful SSVEP information.

Electrode Response To Frequency Range



On screen frequency charted against theta frequency band strength. The red line is a simple regression, indicating correlation between stimulus and response.

As can be seen above, the original 14 electrodes can be reduced to six: AF3, F7, F3, F4, F8, and AF4. A quick glance at the electrode location chart verifies this

reduction, as these electrodes are clustered near the visual cortex. By using experimental data to filter meaningful 'conversations' from noise, Brainstorm is able to feed logistic regression higher quality features which results in higher quality predictions.

Binning

Just because these six electrodes *normally* provide meaningful SSVEP information, doesn't mean they *always* do. The data stream from an electrode can be corrupted by a number of external factors including movement, sensor impedance, and signal transmission. Brainstorm accounts for this noise, by calculating the following frequency power bands over two second bins of neural activity.

Brainstorm Frequency Bands (Hz) — Calculated over 2 second bins and updated twice per second

Theta: 4-8

Alpha: 8-12

Beta Low: 12-18

Beta High: 18-25

Gamma: Greater than 25

By using binned data to calculate frequency bands, Brainstorm boosts each electrode's signal-to-noise ratio; this improves SSVEP classification performance.

Majority Voting

Brainstorm collects predictions over time. For a *final* frequency prediction to be made, the candidate frequency must have received a simple majority of votes. By using majority voting, Brainstorm minimizes the influence of any single prediction which may be corrupted by noise (despite the noise controls detailed above).

4. Output

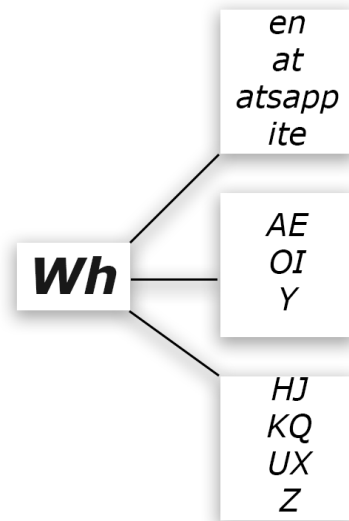
Key Point: *Brainstorm increases typing speed by providing users with autocomplete suggestions.*

After Brainstorm processes EEG data, the predicted frequency is matched to flashing text on the screen. This process is repeated, prediction after prediction, with the final result being a chunk of text. For most virtual keyboard systems, this would be enough. Having added evidence to their thesis, the researchers behind these systems would close their program and call it a day. However, Brainstorm is not an intellectual exercise; it's a practical prototype: a significant step in the long journey to redefine the relationship between humans and computers.

Any system that redefines this relationship needs to be *accurate*, *consistent*, and *fast*. Signal processing and classification take care of accuracy and consistency, while

Brainstorm leverages the NLP technology behind Google to boost *words per minute*, a proxy for typing speed.

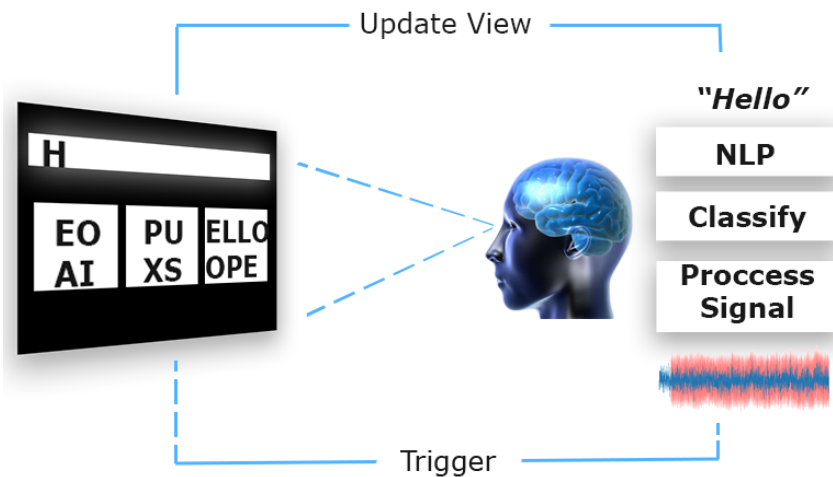
Autocomplete Example



From top to bottom: likely fragments, likely letters and general letters.

Everytime a character is added to the output, Brainstorm queries Google for autocomplete suggestions which are then parsed and displayed to the user. By allowing users to select from high quality suggestions alongside characters, Brainstorm reduces user burden which in turn reduces system burden. The more information Brainstorm can gain from natural language processing, the less information Brainstorm has to extract from brain waves. This leads to an increase in “typing” speed, with Brainstorm achieving an average of 40 characters per minute.

Trigger to Output



The flow from frequency box to brainwave to language.

5. Action

When the user is done typing, the text can be sent to *anyone anywhere* via email. By connecting the user's mind with an established communication platform like Gmail, Brainstorm represents a viable prototype that minimizes the gap between what we think and what we share.

Key Stats

Speed: 40 characters per minute

Code: One open source codebase

Data: An open source data set containing over **50,000,000** labeled data points

Interaction: The world's **first** recorded email sent with a mind controlled keyboard

Future

"A journey of a thousand miles must begin with a single step." — Lao Tzu

The brain is the final frontier. Having conquered Earth, we now turn to conquering ourselves. But as scientists have come to realize, understanding how we think is incredibly difficult. Brainstorm represents a significant step in the long journey to typing with our minds, but there is a lot of room for improvement. Here are a few opportunities for developers looking to extend Brainstorm's capabilities:

- Increase classification accuracy by swapping logistic regression for neural networks
- Adopt a non-SSVEP paradigm that reduces Brainstorm's dependence on external stimuli
- Stress test the software and provide additional error handling where needed
- Add support for EEG systems outside the EMOTIV ecosystem