A cozy computer-human singularity is approaching-- a world where digital interfaces disappear, and people no longer need external devices to communicate over distance. In the future cell phones will not be used to communicate over distance-- people will simply think to send messages to each other. In the future people will communicate with others in a pragmatic equivalent of telepathy.

This future is possible today. In this white report I reveal the reality of the field of telepathy engineering, reveal its basic assumption, and provide direction for AI researchers interested in contributing to this new field. I also describe current progress towards pragmatic telepathy, and a framework "Headtalk" that supports our progress.

The implementation in this report uses multiple sensory inputs including EEGs as wearable devices. This report discusses Coen's multimodal perception framework to yield output symbols from wearable device inputs.

First, what specifically would I like to accomplish? I'd like to use wearable devices to construct a reliable symbolic communication channel between two people. People will be able to assign meaning to the sequences of symbols they send each other, hence building a language. This framework ought to allow the creation of language. My goal with this development iteration is to digitally enable sign-language for people separated by distance.

Second, persistent telepathic communication between people is real-world feasible today. I have used the iPhone as a hub connecting Bluetooth Low Energy wearable devices, processing their data, then transmitting their data to other devices via the internet. An iPhone can run this all day on one charge. In this project, I've built an interface to the EEG sensor from the "Mindflex Duel." This is the first wearable device. The second "wearable" device is accelerometer data from the user's iPhone.

Third. I've made an assumption on why communicating via telepathy is possible, explaining why present AI is substantial for the task. I intend to make human brains do the majority of the computation under this framework. The assumption is, "given a consistent pathway connecting minds together, brains can figure out how to use it to communicate." Instead of developing sophisticated AI that conforms to variations in users thoughts and actions, we have users conform to a consistent AI.

Fourth. What does the AI do? The role of AI is to 1) recognize user intent to send symbols to their partner with high consistency, 2) eliminate false positives, and 3) increase data throughput. I found that Coen's 2006 thesis on Multimodal Perception described a methodology that met all these requirements, and I utilized its methods as a first demonstration. I provide four different research directions towards pragmatic telepathy.

I speak to practical applications. Headtalk is in testing phase with their "Headtalk Magnet" product. Their device is a bracelet allowing Deaf people and their friends to "tap each other," getting attention and transmitting short symbolic messages whether across the room, or across the world.

Contributions described in this report: I,

I) Defined goals for telepathy, explained feasibility
II) Presented mental plasticity assumption
III) Implemented Core-Telepathy framework for iPhone-- Headtalk
IV) Gave four research directions in AI for telepathy

Research was performed within MIT's Aeronautics and Astronautics Dept, and Dept of Electrical Engineering and Computer Science, in collaboration with Headtalk Inc. 222 Third St. Cambridge, MA 02142.

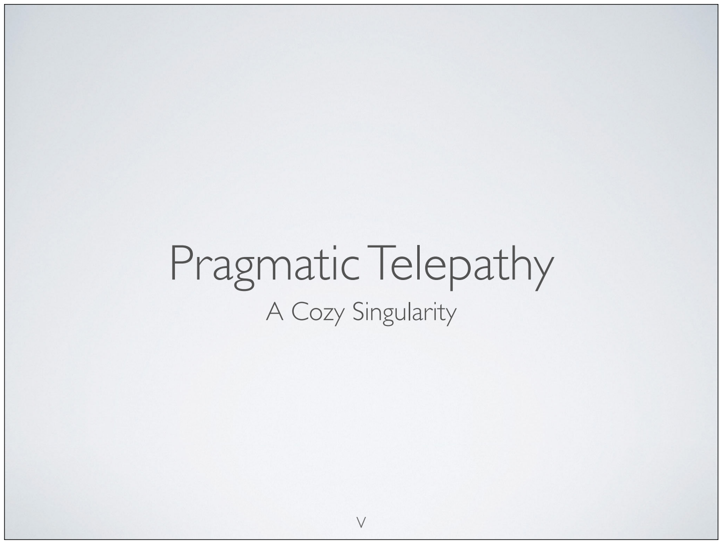

In the future cell phones will not be used to communicate over distance-- people will simply think to send messages to each other. In the future people will communicate with others in a pragmatic equivalent of telepathy.

Alexander H List
Research performed within MIT Aeronautics and Astronautics Dept, and Dept of Electrical Engineering and Computer Science, in collaboration with Headtalk Inc. 222 Third St. Cambridge, MA 02142.

Special thanks to Professor Patrick Henry Winston, advisor, mentor and professor. Ford Professor of Electrical Engineering and Computer science at MIT CSAIL laboratory.

If we are to have a cozy singularity, we ought to use people's own bodies as an input device. We better get rid of external sensors, get phones out of the way, and remove any requirement that they do one thing at a time. What's more, communications better be intentional and robust.

As the simplest complex example, imagine two people who are able to send morse-code-like messages between each other, just by tapping on their wrist.

- I'd like to use wearable devices to construct a reliable symbolic communication channel between two people. My goal with the discussed implementation is to enable digital sign-language for people separated by distance.
- People will be able to assign meaning to the sequences of symbols they send each other, hence building a language. This framework ought to allow the creation of language.

The steps this paper takes are 1) describe goals, 2) explain feasibility, 3) present the "Mental Plasticity Assumption," and 4) describe the role AI frameworks have in making telepathic communication possible, specifically discuss Coen’s multimodal perception framework. These will be followed by brief discussion in the form of news, and contributions.

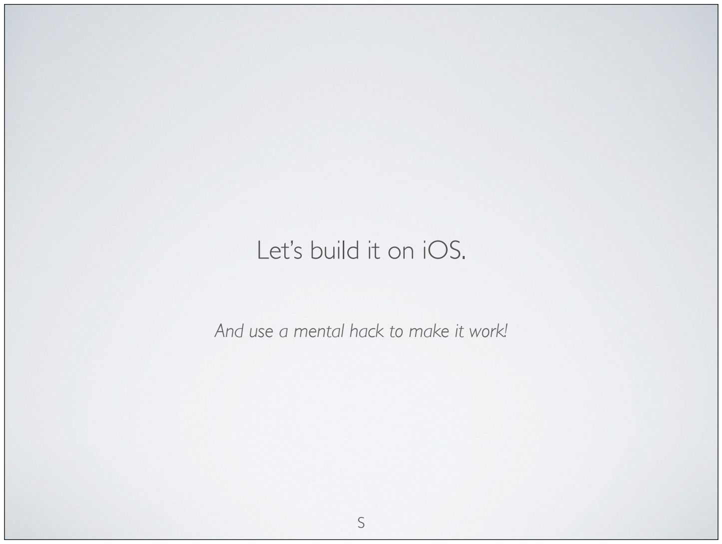

This is all possible using the iPhone as a hub.

With prevalence of iPhone, instantaneous persistent networking, and Wearable Bluetooth Low Energy devices, all technologies are for the first time in line for pragmatic telepathy. Low Energy Bluetooth for Wearable Devices lets data acquisition always be running. Plus, almost any sort of physiological data is available from off-the-shelf wearable devices.
Recent developments in AI like Coen's multimodal framework seem viable for processing their inputs, and run quickly within low power requirements on the iPhone in C++. Persistent internet sockets are essentially free for connecting phones together, so people can communicate instantly.

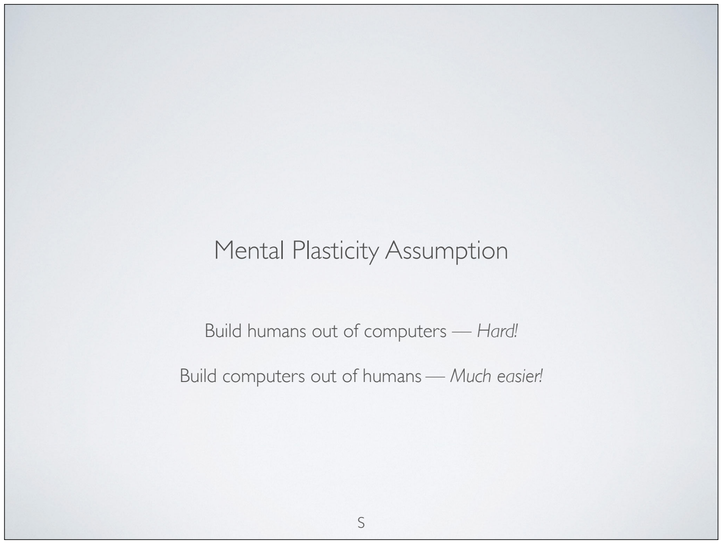

The assumption is, let's work hard to make the AI consistent, and let the brains figure out how to use it!

I intend to make humans do the heavy computation in this framework. Not the computers. I commonly contrast the try-hard uselessness of Siri, and the way call centers have become more automated.

I have the realization that trying to build a human-like creature from machines is a time-wise asymptotic endeavor. Building a machine intelligence starting from networks of humans is similarly feasible, yet much more economically valuable.

Instead of developing sophisticated AI that conforms to variations in users thoughts and actions, we have users conform to a consistent AI. I employ relatively straightforward multimodal AI that coalesces input data, and provides consistent outputs. Hence "users can get used to it" and use the pathway to "create their own languages."

This seems rather reasonable. People are used to communicating on restricted channels: morse code and sign language exist. Humans invent languages on the fly-- two children create their own language to play with each other, "one clap means jump!" Spouses lexify so their children can't understand them, "so-and-so made a code 52 in the shower."

Engineering is simpler when we do not need AI from the future, so we put our focus into making the AI consistent rather than flexible.

In this project, I've built an interface to the EEG sensor from the "Mindflex Duel." This is the first wearable device. It samples at a maximum rate of one sample per second, and gives time averaged concentration values from overall brain activity.

The second "wearable" device is accelerometer data from the user's iPhone.

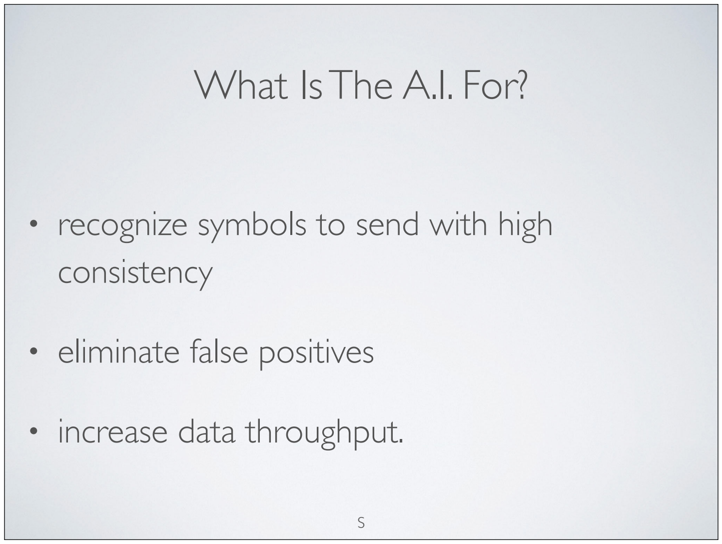

Compared to external devices, the trouble with wearable device-based control is that they are prone to accidental activation. People wear their devices all day. While video games utilizing EEGs and accelerometers can demand total constant attention, devices on standby for telepathy can not.

If the probability of an accidental activation on each mode is Pa, requiring two modes to activate to send a symbol makes the probability of accidental activation Pa^2. Generally, the system relies upon a combination of multiple simultaneous intentional actions to ensure sent symbols are intentional.

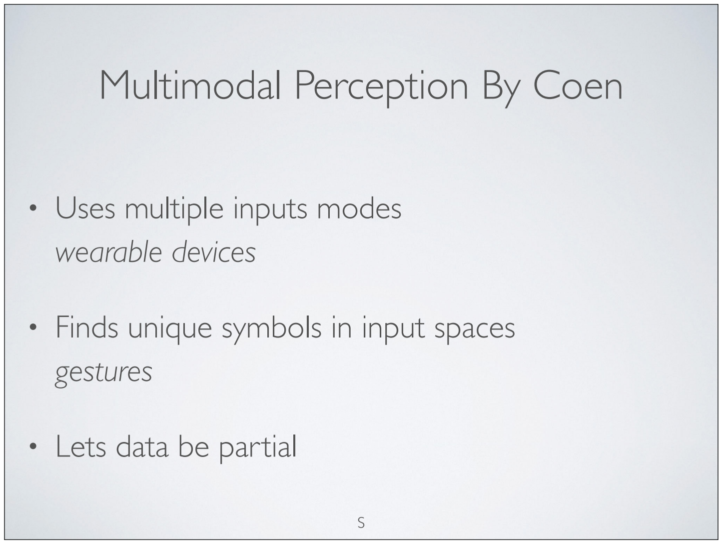

Coen's 2006 thesis on Multimodal Perception described a methodology that met all these requirements, and I utilized it for the purposes of first demonstration.

Coen's framework requires a bootstrapping phase where users perform gestures they'd like to be independently recognized. This data is processed using multimodal clustering algorithms that separate inputs from each wearable devices into symbols.

The value of this algorithm is that it tends to extract a larger number of potential output symbols from an input dataset. It conducts unsupervised learning on a training data set.

The unsupervised learning process has the side-effect of allowing similar gestures, such as where their difference is an EEG reporting a user's high versus medium-high concentration level, to be symbolically equivalent.

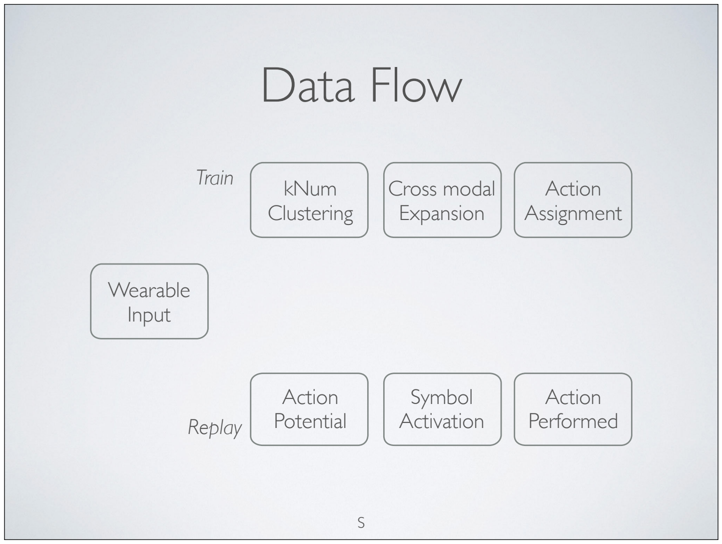

I developed the AI module in Headtalk to be drop-in replaceable. Coen's framework required a train mode and a replay mode. Data flow from left to right in both cases.

I have implemented the majority *train* data flow on iOS using a mix between C++ and Objective-C, was +/- 3500 lines of code including Headtalk cross device telepathy code.
Data is stored on device between app runs using a sql-lite database.
Coen's piece was roughly 1000 lines of code that I implemented, I wrote kNum clustering to get specific numbers of points per cluster.

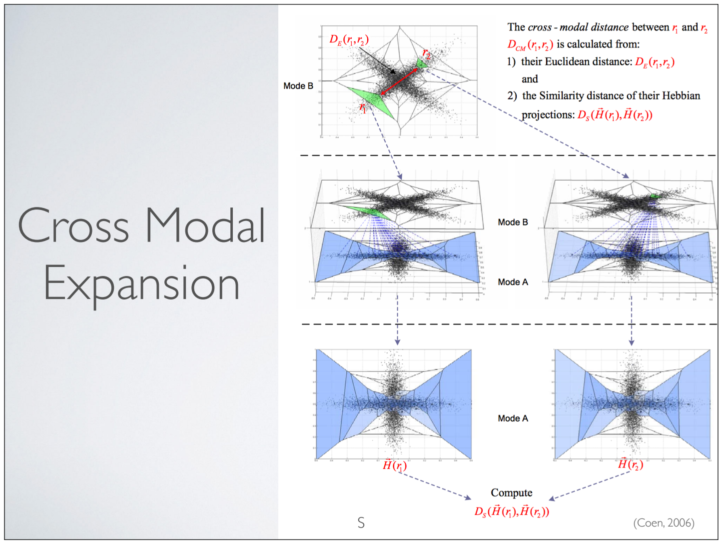

From Coen’s thesis, page 74
Suppose mode B is the the accelerometer x,y input plane.
Suppose mode A is the input plane of two EEG diodes reading concentration levels.

Two clusters in Mode B, are being tested to see whether they belong in the same symbol so they can be merged.
In Coen’s framework, the test is performed by observing “the way each cluster corresponds to data in other modes.” For each cluster, the question is, what other data points occurred at the same time as data points inside my cluster? Hence a probability vector is created from the test cluster in B to each cluster within A. Elements for the probability vector of the first test cluster in B are [# of same-time corresponding points from cluster B1 to cluster Ai / # of total points within B1].

Probability maps for the first and second cluster are compared with a distance metric. While Coen describes a nondimensionalized distance that removes need for a minimum closeness parameter through principles of averaged disorder, I have not spent substantial time to corroborate the functionality of this metric.

Professor Patrick Henry Winston instead suggests using a cosign distance between probability of each cluster, and applying a minimum closeness parameter. This technique is far more computationally efficient, and works especially well when data is well distributed. Cosign distance does have difficulties. Cos is susceptible to “dimensional bucketing,” where in off-dimension clusters in vector one have no effect on closeness corresponding to a specific cluster in vector two. Namely, cosign doesn’t appreciate neighborhoods or proximity of clusters.

Nonetheless, I implement the cos metric for it’s computational and conceptual simplicity. When cosign, the dot product divided by magnitudes squared of the test clusters’ probability vectors, is greater than 0.7, they are merged into one symbol.

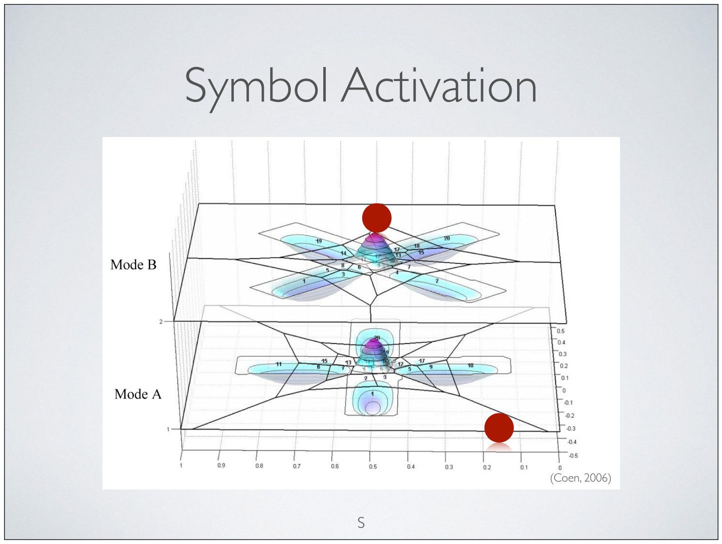

Coen's framework limits false positives by having action potentials for every symbol. In "run mode," the closer input data from wearable devices is to original gestures, the faster they activate. If the input data is different, then it is slower to activate. Hence single-device inputs won't accidental cause activations as readily.

Above, red dots on mode A and B represent “activation cursers.” Cursers are moved closer to a specific symbol when their device’s input corresponds to that specific symbol’s cluster, OR when another curser for another input corresponds to that same symbol.

From Coen’s thesis, page 101, partial data is also managed by having “repellers” and “attractors.”
“Repellers, corresponding to ambiguous nodes, are maxima and shown in the center; attractors, corresponding to unambiguous nodes, are minima and are outlined by their basins of attraction..”

The final role of the AI is to increase the data throughput of communication from input data. The way Coen's AI does that is by increasing the quantity of symbols available for activation in self-supervised learning. An alternative way is through increasing the speed of symbol activation. We can tune Coen AI's potential activation strength parameter to change recognition speed, though this way is inherently more prone to false positives.

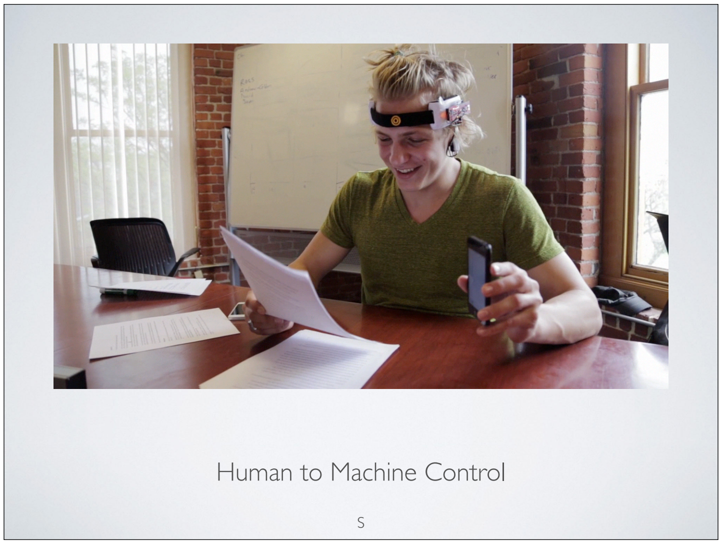

video http://youtu.be/mZIr2XDfgIE

As of this writing, only a naive implementation of Coen's framework has been achieved. While the train flow should be implemented by June 2014, the input devices used have still been extraordinarily crude. For example,

It is not yet conclusive whether are no unknown unknowns that AI needs to account for.

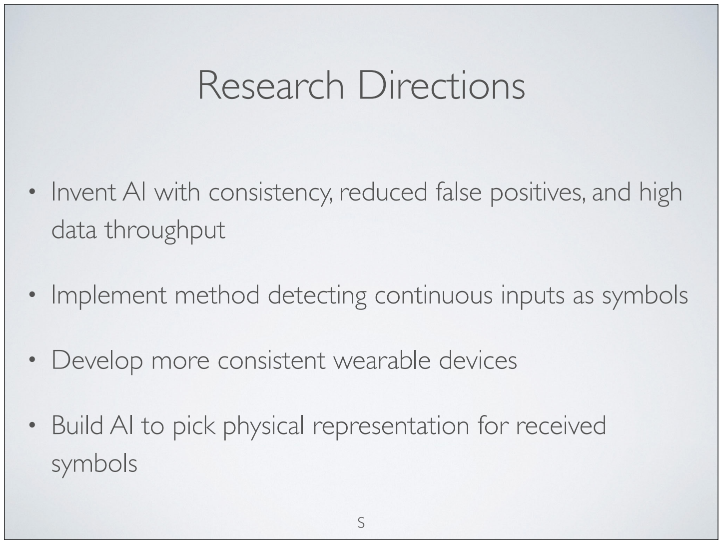

The general direction for research into enabling telepathy should be in simplifying AI implementations to achieve goals of 1) consistency, 2) reduced false positives, and 3) high data throughput.

Further investigation should be done into perceiving continuous inputs into symbols. While Morse Code is composed of discrete symbols, American Sign Language has symbols that are recounted as continuous time actions. McIntyre's 2014 masters thesis on Embodied Representations could hold some clues.

Research also ought to be done from the device perspective. Accelerometers return data relative to earth's gravitational field. Devices that present acceleration relative to a user's body orientation would be more effective at providing intelligible and consistent data to perception frameworks.

AI also participates in choosing a physical representation for symbols received from a partner. For a receive symbol a tone can be played, lights flashed, or vibration activated. There is much work to be done in this direction to determine how best to articulate a perceivable phenomena for a received symbol. In this implementation audible clicks were heard by recipients. There is much research to be done in this field.

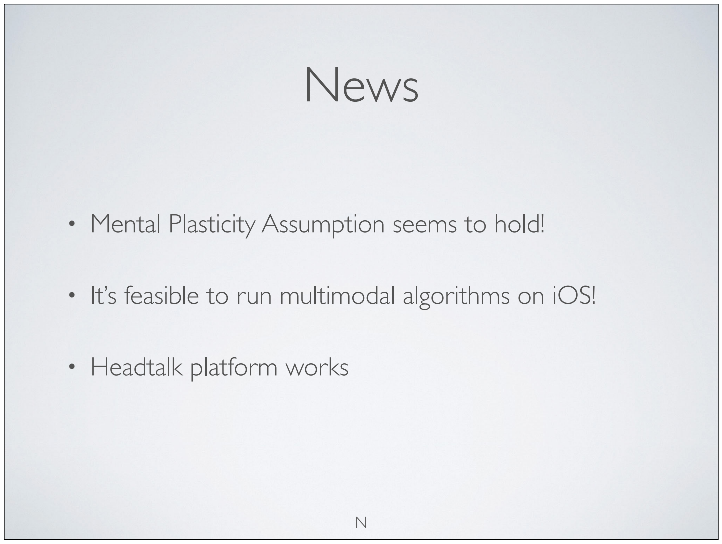

Research collaborator Headtalk Inc has compiled footage that corroborates the Mental Plasticity Assumption, has shown their system with multiple wearable devices in action.

Interestingly, running algorithms on iPhone 5 is only about 20x slower than running on high performance laptop computer. The current algorithm in O(N^2), but there is hope to refactor to an O(N log(N)) algorithm.

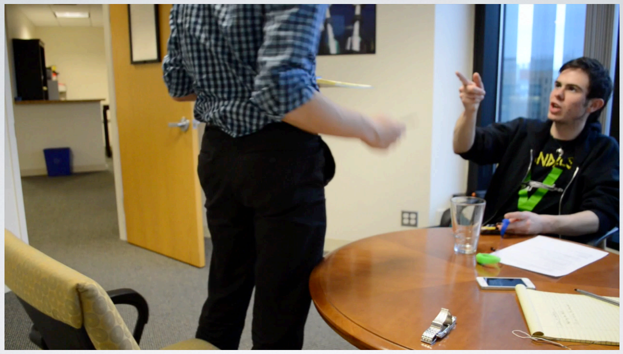

Humans Communicate

N

video http://youtu.be/-BLG7BvXZ0c

Showing a test where Right is asked to share what genre the music is to Left by tapping on their bracelet. They first spend a few minutes to design their own language, then they are able to successfully convey tempo and genre information to each other.

While it does not validate, it certainly corroborates the Mental Plasticity Assumption.

Testing A Single-Mode Device

For Deaf People, Athletes, and Children

"Magnet" prototype right

N

Headtalk is in testing phase with their "Headtalk Magnet" product. Their device is a bracelet allowing Deaf persons and their friends to "tap each other," getting attention and transmitting short symbolic messages whether across the room, or across the world.

Headtalk tap bracelet utilizes one input stream. While its false positives rates are low, its data throughput rates are low too.

Alexander H List, ©2014
alist@mit.edu