

Moran Cerf, Hagar Gelbard-Sagiv, and Itzhak Fried

Much of cognitive neuroscience, investigated in humans and animals, is based on careful evaluation of neural responses to sensory stimulation or external cues. This is the ultimate product of the stimulus–response paradigm, which has been the cornerstone of neuroscience research. Yet, a considerable extent of mental life is spent in thought processes that may not arise as a direct response to external input. Such processes, which include imagery, free recall, and internal deliberations, are the very processes which are difficult to study in animal models and require the cooperation of human subjects who can express their thoughts and wishes and make them accessible to an external observer.

A basic tenet of modern cognitive neuroscience is that thoughts correspond to patterns of neuronal activity. Indeed, the idea that thoughts may be decoded from such patterns has posed a considerable challenge. It is not that we necessarily want to become “mind readers,” but the prospect of such decoding presents theoretical and clinical benefits, not the least of which is the construction of neuroprosthetic devices that will enable patients that are unable to speak or move to communicate.

While numerous methods of neuroscience research, such as electroencephalography (EEG), functional magnetic resonance imaging (fMRI), positron emission tomography, magnetoencephalography, and others, offer some ways of inferring thoughts from patterns of neural activity (Kamitani & Tong, 2005; Haynes & Rees, 2006; Norman et al., 2006; Kay et al., 2008; Nishimoto et al., 2011), these methods are limited by the nature of the signals measured, including their spatial and temporal resolution of readouts from the brain. It seems that the decoding of discrete thoughts and concepts requires access to a different kind of signal—spikes of single neurons. Such recordings are feasible in humans under special clinical circumstances, thus enabling investigation of signals at the neuronal level in human subjects who are able to report their imagery, recollections, and other mental experiences directly and who are often able to elicit such processes at will when instructed to in an experiment.

In a series of studies conducted recently using single neuron recordings in humans, scientists were able to decode internal mental processes, in relatively high specificity, and in some cases prior to their actual manifestation by language or motor behavior. These pioneering studies form

the scientific basis for further investigation into the possibility of decoding freely occurring thoughts and deliberations.

Looking inside Our Mind's Eye Using Imagery

Trying to direct another person to our home, or trying to recall a detail from a picture we have seen before, we often close our eyes and reconstruct an image in our mind's eye. This employs an internal process by which we visualize content that is not present on our retina, using imagery. The single neuron correlate of this internal process was studied by Kreiman and colleagues (Kreiman et al., 2000) by looking at medial temporal lobe (MTL) neurons in patients with intracranial depth electrodes.

In a preliminary stage subjects were shown a variety of images from a selected set of categories: objects, familiar faces, animal images, and so forth for 1 s each in the course of a short viewing period. Following this viewing period, the activity of a small number of brain cells was analyzed, and a set of neurons were identified, each showing a clear response to images from a particular category (e.g., a neuron showing a response to animals and another showing a response to objects). These neurons were then used for a subsequent experiment where subjects were asked to imagine the images seen previously. Subjects were asked to first view two images from two categories, shown repeatedly in a random order, for 1 s each, with an accompanying simultaneous tone. The tone was either a high-pitched or a low-pitched one. The subjects quickly learned to associate the high-pitched tone, for example, with the image of the animal, and the low-pitched tone with that of the object. These sets of trials were labeled "vision trials." Next, "imagery trials" followed. In these trials subjects were asked to listen to the high/low-pitched tones while closing their eyes and imagine the picture corresponding to the particular tone. The tone therefore triggered the internal imagery in the subjects' minds. The tones alternated in a sequence allowing the subject to perform alternate imagery of the two different stimuli. Sixteen percent of the recorded MTL neurons showed selectivity during vision—that is, they reliably responded with an elevated firing rate each time a specific image was presented and remained silent while other images were presented. Eight percent of the recorded neurons showed selectivity during imagery. Remarkably, almost 90% of the imagery neurons also showed the same selectivity during vision.

Comparisons of the neuronal response characteristics between vision and imagery in this experiment show various differences in latency and duration, as expected, taking longer for images to imprint via our imagination than through vision (see figure 11.1). Yet, two striking findings emerged. First was the presence of a set of neurons responding selectively to one image, which then showed the same selective responses during imagery. The second was that the firing rate for imagery was 85% that of vision. This suggests an only slightly diminished response in our mind to imagery with no retinal input relative to normal vision. Ultimately, these results enable an external observer an access to the internal imagery process. One could effectively parse the stimulus imagined by the subject simply by the readout of the activity of a few corresponding




Figure 11.1 Individual responses of a single neuron during vision. Two images were shown separately for 1 s each, with five repetitions per image, indicated by horizontal black and gray bars. After each picture, subjects pressed a button to indicate whether or not the picture was a human face. The continuous thin black line shows the spike density function. After ten visual presentations, subjects closed their eyes and imagined one picture on hearing a high tone and the other picture on hearing a low tone. Tones were alternated every 3 s. Data shown are from the same neuron during visual imagery. This neuron showed a similar pattern of firing during visual presentation and visual imagery, with increased firing rate in response to the image of the ball, but not that of a human face (adapted from Rees et al., 2002).

individual neurons. Indeed, on average, in seven out of ten trials, an observer could correctly guess the stimulus category in imagery or visual trials simply by counting the number of spikes of the neurons selective to these categories. Presumably, given a larger set of neurons from the population that encodes a particular category, the prediction rate could be much higher.

Thoughts and Recollections

Recollection is another form of thought that is often triggered in our brain. Retrieving information from the past requires an internal search process that ultimately results in the recall of a certain event or a piece of knowledge. Several studies have tried to explore how similar the neuronal activity patterns during recall are to the ones evoked during the actual original event (Nyberg et al., 2000; Wheeler et al., 2000; Polyn et al., 2005). However, given the limited resolution of imaging methods, and the difficulty in accessing mental content in animals, this question has remained open.

Gelbard-Sagiv and colleagues (Gelbard-Sagiv et al., 2008) directly investigated the underlying neuronal correlates of recollection by having human subjects freely recall previously shown content while recording single neurons from their brains.

Thirteen subjects first viewed a series of short (5–10 s long) video clips. Each clip was repeated six times in random order and depicted an episode featuring famous people, characters

or animals engaged in activity, or landmarks which were often familiar to the subjects. These clips included well-known scenes, for example, the famous scene of Marilyn Monroe's dress floating around her above a subway grate, or President Bush's speech after Saddam Hussein's capture. About half of the recorded neurons responded reliably and selectively during one or more of the clips. Ten percent of all responsive neurons showed sustained responses, that is, an elevated firing rate that was maintained throughout the clips, and also after the clip removal, when no content was presented. In some cases this prolonged firing was attenuated only by the onset of the following clip. Following the viewing session, and after performing a distracting task, subjects were asked to freely recall as many of the previously seen clips as possible and to verbally report immediately whenever a specific clip came to mind. Thus, the retrieval of information was triggered internally in the patient's mind and not by an external cue. Similar to the viewing session, this free-recall session was designed as a simple and natural task in order to mimic real-life recollections. Indeed, subjects performed well and, on average, recalled over 80% of the clips presented. Interestingly, cells maintained the same selectivity in the viewing and in the free-recall session. That is, a cell that was firing during each of the six presentations of the Marilyn Monroe clip fired again just before the subject reported recalling that particular clip. This result extends our understanding of the abstract nature of MTL representations: Not only do neurons in these regions respond to very different external world representations of a given episode but they also respond when retrieving the content internally. In the imagery experiment we learned that a simple tone associated with a given concept could trigger the same neuronal response during imagery whereas this study demonstrated that no external cue was required—simply “thinking of” or recalling a concept is associated with the same neuronal firing patterns evoked by sensory exposure to it (see figure 11.2 for an example). Remarkably, neurons started firing in the free-recall session 500–1500 ms before the onset of verbal report of the recall of their preferred clip. Although we cannot completely rule out the possibility that this was merely an issue of reaction time or that the subjects were waiting for a second or two before reporting their thoughts, the early neuronal activity observed here could represent the early stages of the formation of a thought or recollection, the stages at which we may not yet be fully aware of the content of the thought. Thus, it is possible that by monitoring the neuronal firing online, we could have told what the subjects were recalling or thinking before they were able to do so.

Decoding a Thought from Neural Activity

Ultimately, our experiences, our knowledge, and our thoughts are represented as patterns of neuronal activity in our brain. Thus, decoding a thought is analogous to figuring out the corresponding pattern of activity that underlies a particular thought.

In our recordings we can only sample a very small subset of the entire population of neurons. However, given the sparse and thus nearly binary nature of the MTL neuronal responses (i.e., a strong bursting response for a specific concept and low baseline firing during other times), decoding a particular concept from this small subset is nevertheless feasible.

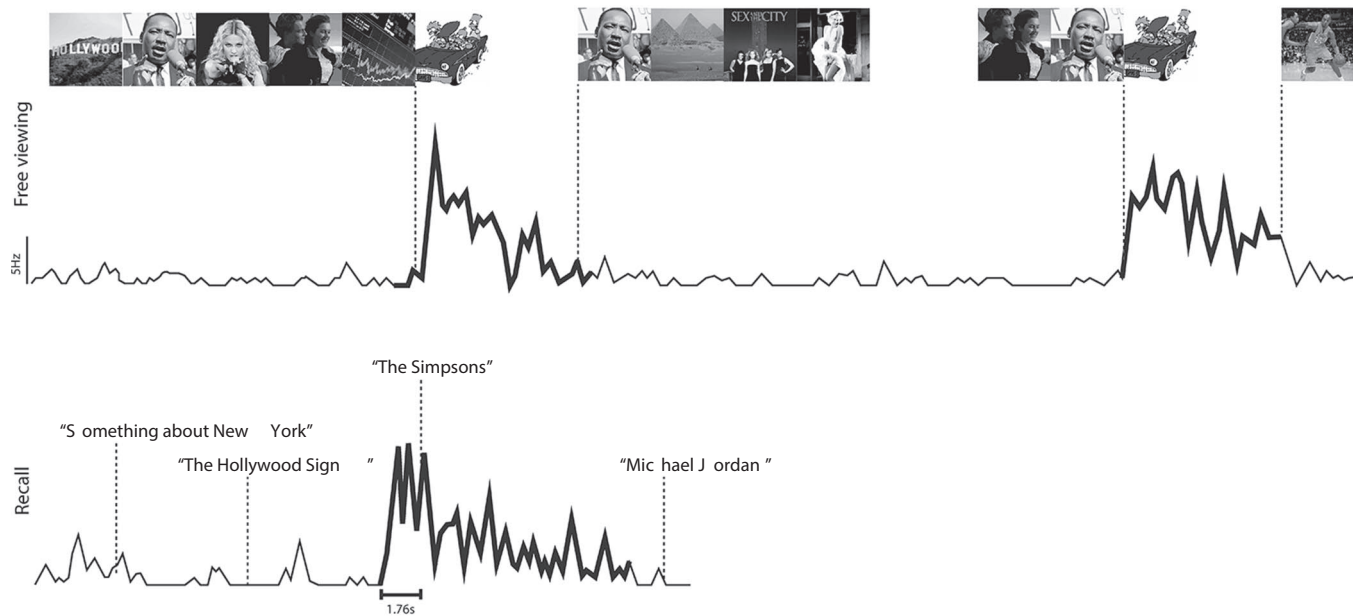


Figure 11.2

Subject freely views video clips (top panel). Each 5-s clip is depicted here by one frame; 5-s blank periods appear in between some clips. Note that the response is sustained throughout the blank period even though the clip is no longer on the screen. Subject freely recalls the clips (bottom panel). The increased firing rate response indicating the recollection of the “Simpsons” clip starts over 1 s prior to the verbal expression of that recollection.

In a recent study by Quian Quiroga and colleagues (2007), neuronal data collected while 11 patients were watching a set of images, presented for 1 s and repeated six times each, was used to quantitatively test our ability to decode the stimulus presented from the patterns of neuronal firing. Baseline firing was defined as the median firing rate from 1000 to 300 ms before stimulus onset. Similarly, the response to a specific image was defined as the median firing rate during the 300- to 1000-ms period following image onset. Neurons that elevated their firing rate to more than 5 standard deviations above the baseline firing rate for a given image were considered responsive to that image.

Using a leave-one-out decoder (using the five other image presentations to “train” the decoder), Quian Quiroga and colleagues could decode significantly above chance which picture was presented to the subject on a trial-by-trial basis. Decoding performance increased linearly with the number of units used as an input for the decoder. In accordance with the invariant nature of the MTL representations, decoding was not limited to repetition of a single photo exemplar but could also be generalized to new pictures of a given person or landmark.

Figure 11.3 depicts the decoding performance in one of the sessions in the form of a confusion matrix. Based on the activity of 19 responsive neurons, the investigators could decode which

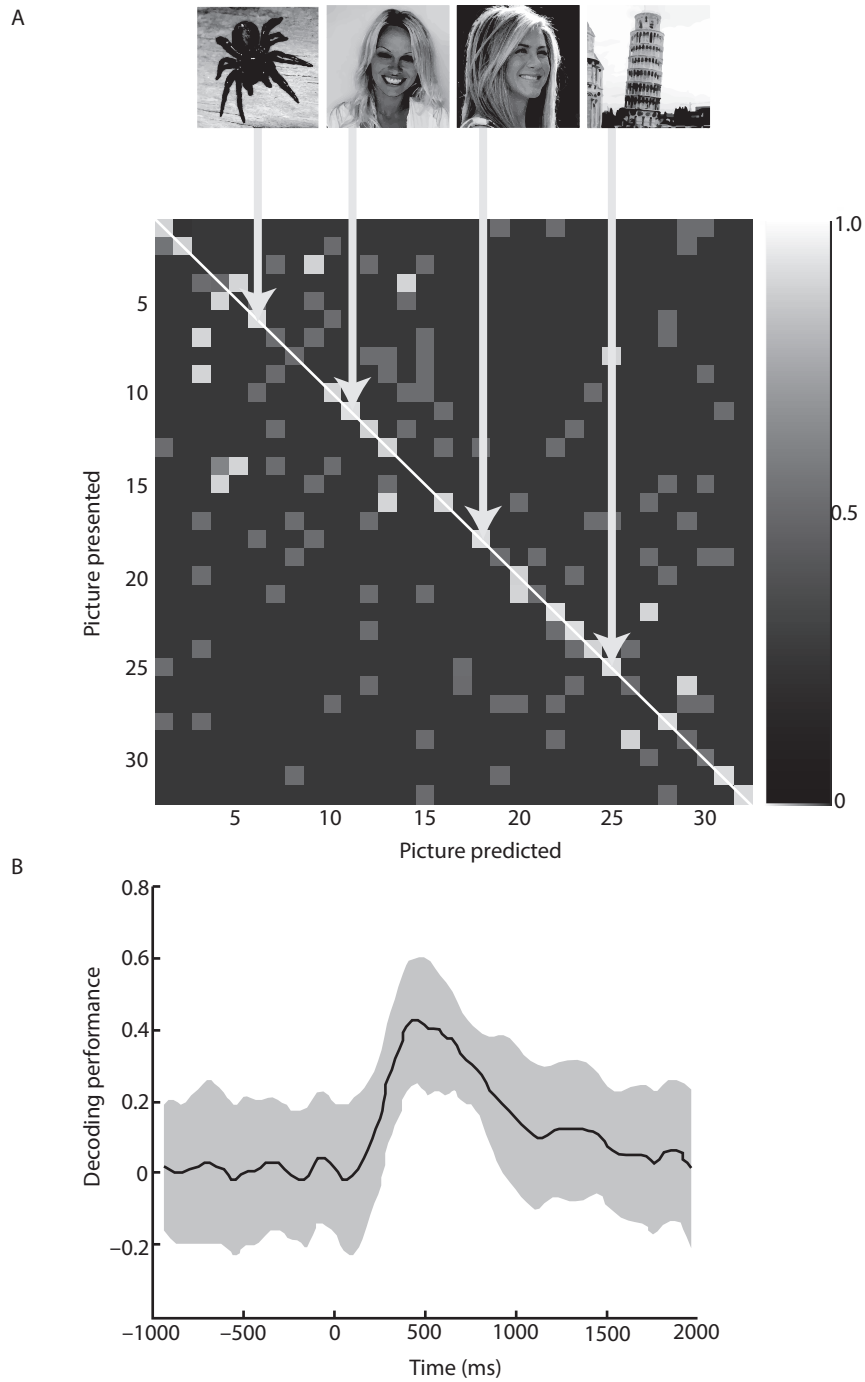


Figure 11.3
 (A) Example decoding of 32 pictures that generated a significant response in any of 19 responsive units in one session. Average decoding performance for all 32 pictures was 35.4%, with a chance level of $1/32 = 3.1\%$. (B) Time profile of the normalized decoding performance using a (half-overlapping) moving window of 100 ms, averaged across all sessions. Band shows the 95% confidence intervals. Performance peaks between 400 and 500 ms. Values were smoothed using a 3-point moving average (adapted from Quiroga et al., 2007).

of the 32 different pictures was presented onscreen in any given moment with a probability of 35.4% (compared to 3.1% chance level).

Overall, the decoding performance first became significant 300 ms after stimulus onset with a 20% decoding performance; it then peaked between 400 and 500 ms with a performance of around 40%, followed by a gradual decay (see figure 11.3). The time window between 300 and 600 ms contained the most informative spikes, and the decoding performance in longer time windows was not significantly higher. Additionally, it was shown that using isolated putative single units resulted in superior performance relative to using multiunits (without spike sorting).

This study demonstrated the feasibility of decoding a visual stimulus presented to the subject from the activity of a few responsive neurons in the human MTL, ultimately showing that the performance relies mainly on an average of 4–5 spikes occurring between 300 and 600 ms after image onset.

Reading Thoughts in Real Time

Following the work by Quiroz Quiroga and colleagues, Cerf and colleagues (Cerf et al., 2010) tested the ability to use population decoding to identify a subject's thought in real time. Using a similar method, patients were first presented with a set of four images that elicited neuronal responses in an earlier screening session. Images were presented for 1 s, 12 times each, in random order. Based on the neuronal firing during image presentation, one target neuron was selected for each image such that this neuron responded strongly to that image and was indifferent to the other three images (i.e., remained at baseline firing when the other three images were presented). The firing rates of these four target neurons in the 300- to 700-ms interval after stimulus onset were represented in a four-number vector. This yielded a single data point in a four-dimensional (4D) space for each image presentation. The 12 data points for each image (12 repetitions of that image) create a cluster in this 4D space. Thus, when the responses are strong and reliable, the data points for the four images create four clearly separated exclusive clusters.

Based on these four predetermined clusters, Cerf tested the ability to decode the internal representation of the patient in real time. The subject was asked to voluntarily think of one of the four preselected images. The activity of the four target neurons during 100-ms time windows was projected to the 4D concept space set before, and using a distance metric from the center of each concept centroid, each momentary thought was assigned to the cluster to which it belonged. The algorithm could also assign a "none" result, meaning that the new data point could not be identified as part of any of the clusters. Based on the result of this decoder, an image corresponding to a concept cluster gradually appeared ("faded in") onscreen. The contrast of the image was increased as neuronal activity was repeatedly identified as part of the same cluster in 100-ms intervals. This served as a real-time feedback for the subject, who could control the visibility of an image by focusing his or her thoughts on a concept. This reflects a gradual transition from a thought or imagery limited to the subject's brain alone to a visual manifestation of a subject's thought.

Next, Cerf and colleagues targeted the ability to override visual input and control the competition between multiple thoughts/concepts in the brain. Similar to the “fading-in” experiment described above, subjects were instructed to make a target image more visible; however, this time a superposition of two images (the target image and a competing image) was presented on the screen. Each of the images was the preferred image of a different neuron in the set of four neurons used by the decoder. The decoder result was used to update the relative contrast of each of the two images superimposed on the screen. The starting point was 50% contrast for both images. The subject then had to not only increase the firing of the neuron (or neuronal network) coding the target image but also suppress the activity of the neuron coding the competing image while the competing image was visible onscreen. Subjects’ performance on this complex task was surprisingly high, with over 70% success by the eighth trial (see figure 11.4, plate 11).

Activating External Devices Using Thoughts


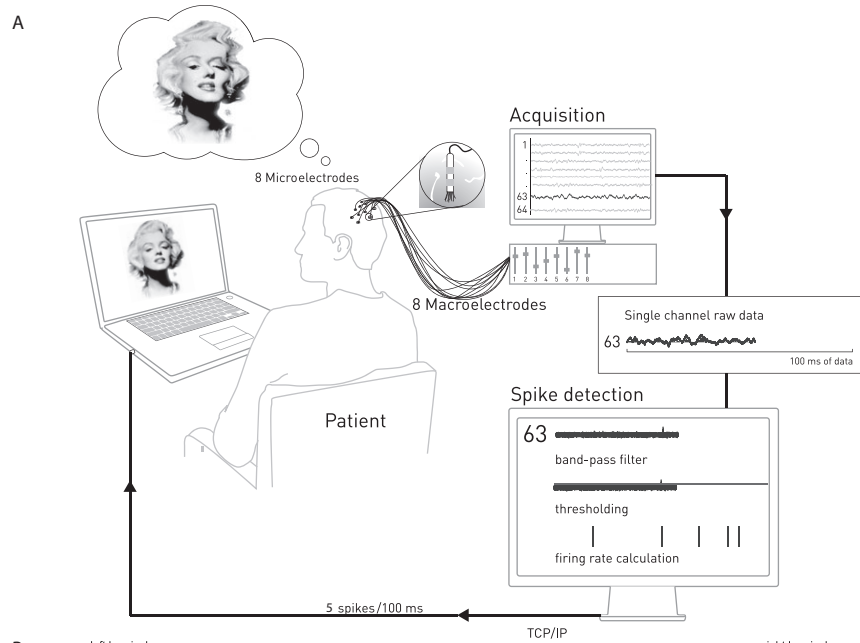
The ability to decode one’s thought with such a high level of accuracy has implications beyond the pure understanding of the ways by which concepts are encoded in the MTL. One obvious implication is the precise control of an external device using these signals.

To test for the feasibility of this, Cerf and colleagues conducted another experiment in which subjects selected the image that they were able to fade in on the screen in the previous experiment most easily and were playing a computer game where the activity of the corresponding neuron was used to control a spaceship (see figure 11.5). The spaceship was moving while obstacles were appearing on the screen—heading toward the spaceship. The subject’s task was to try to avoid these obstacles by controlling the height at which the spaceship flies. He or she could do this by thinking of/imagining that particular image and thus increase the firing rate of the corresponding neuron and make the spaceship fly higher or, by “suppressing their thinking of that image,” decrease the firing rate of the corresponding neuron and make the spaceship fly lower. Subjects were able to control the activity of an individual neuron (or a neuronal network of which this neuron was a part) to the level of roughly 70% performance on that task. When debriefed, subjects mentioned that they mainly tried to imagine the target image.

Similar to the control of the spaceship in the example above, decoding of signals from the brain with high precision can be used to control devices such as a robotic arm (Velliste et al.,

Figure 11.4 (plate 11)

(a) Recording from intracranial electrodes, neurons are identified that respond to a specific concept: in this instance, a cell responsive to the image of Marilyn Monroe. This cell increases its firing rate to the image or thought of Monroe. (b) This cell is then pitted against a different cell which was found to represent the Eiffel Tower. The two images are superimposed, and the subject is asked to bring the image of Monroe to maximum visibility. The visibility of the image is controlled by real-time decoding of the activity of each neuron relative to the other neuron and its own baseline. In this example, we show a case where the subject initially begins to fail the experiment—the firing of the Eiffel Tower neuron increases and the visibility of the tower increases, creating negative feedback. However, the subject is able to exert control and, by concentrating on the internal thought of Monroe, is able to override this sensory input and increase the firing rate of the Monroe neuron and decrease that of the Eiffel Tower neuron, bringing the image of Monroe to visibility. The scans show the location of the respective electrodes within the brain.



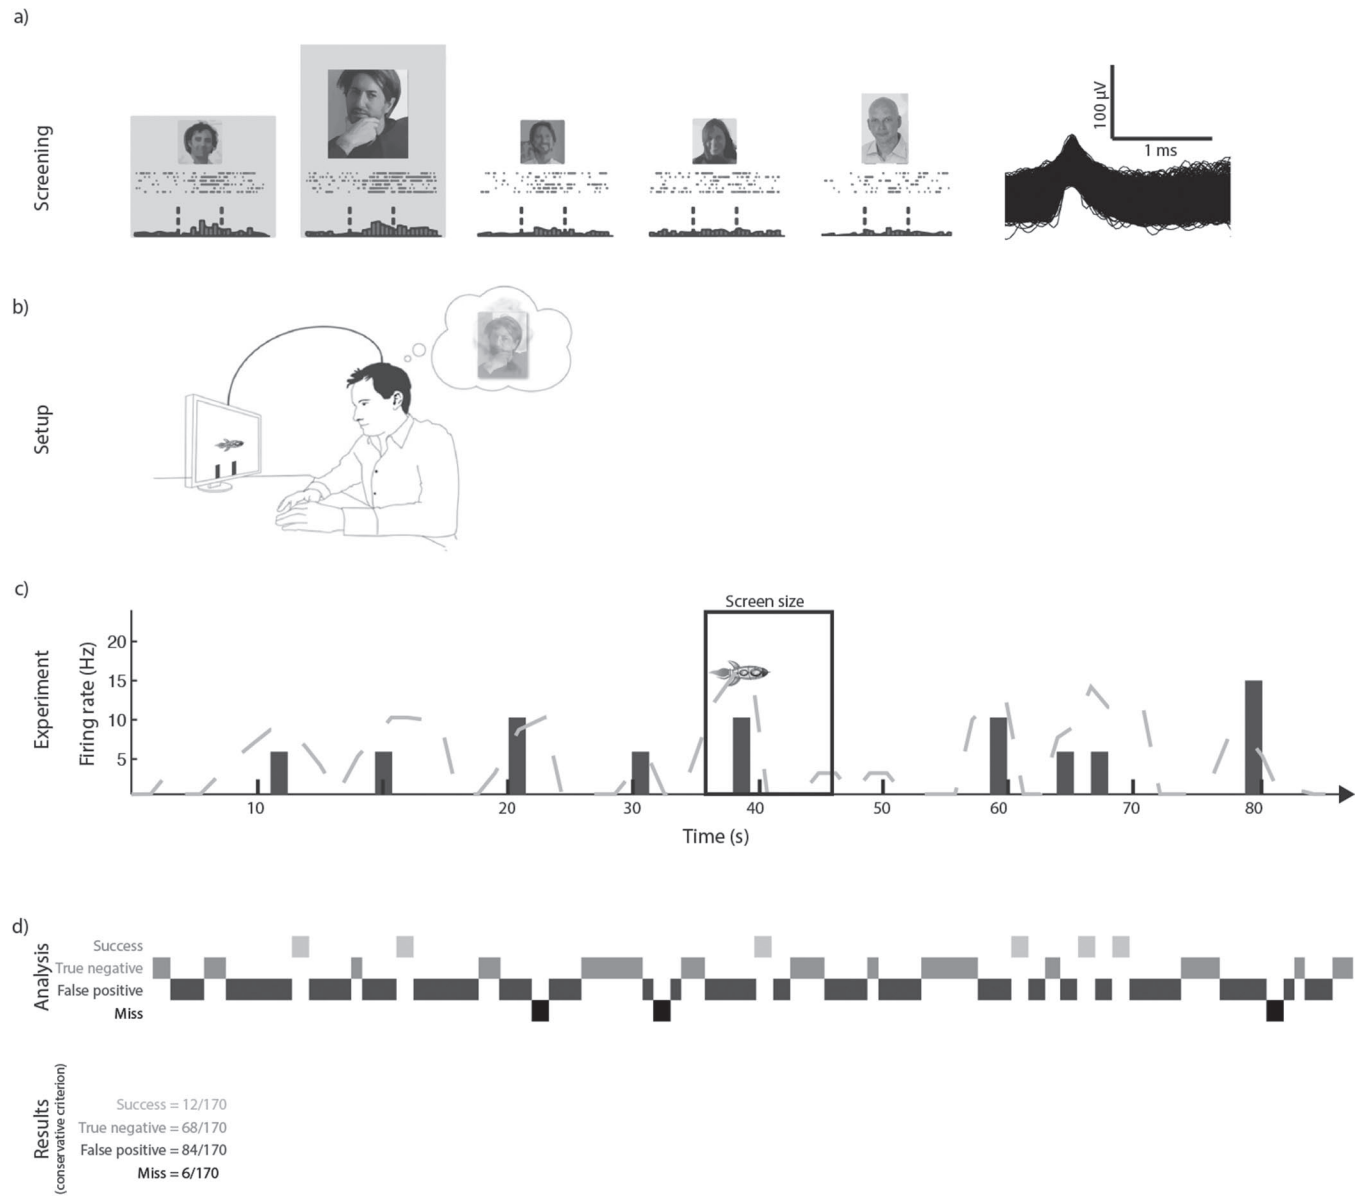


Figure 11.5

(a) Poststimulus time histogram of a single unit in the left hippocampus that responded selectively to a picture of the first author and various other people the subject was familiar with. Each image was repeated 16 times during a screening session prior to the brain-machine interface experiment. To the right are the spikes shapes. The two images with strongest response are highlighted, with that of the first author, which was used for further experiments, magnified. (b) The experimental setup. The subject was watching the game on the screen and controlled the flight of a spaceship by thinking of the first author. (c) A depiction of 90 s of the game. The black bars are the obstacles the subject was asked to avoid by elevating the firing rate of the hippocampal unit and, accordingly, the spaceship's height. The rectangle marked "Screen size" shows the part of the obstacle course the subject sees at a given window (7 s ahead and 3 s before). (d) In order to evaluate the performance, the subject was scored by the amount of hits/misses and true/false positives. A hit was considered when the spaceship hovered above an obstacle, a miss was colliding with one, a true positive required being on the ground when no obstacle is there, and a false positive involved hovering above ground when not needed. The scores match the route shown in (b). Below are the total scores based on the route shown.

2008; Hochberg et al., 2012). In this case, the decoding of neural activity serves as a brain-machine interface (BMI).

The advantage of using MTL neuronal activity for purposes of BMIs is in the accuracy of the decoding, which can outperform other methods such as external noninvasive methods to field potentials (Lebedev & Nicolelis, 2006). Spikes from a neuron that indicates a specific concept are highly informative and thus have a high signal-to-noise ratio. Additionally, the latencies of these neurons are quite short and allow for rapid control. As the free-recall study demonstrated, neurons start increasing their firing rate well before subjects reported their recollection. Given these properties, BMI that utilizes these recording can be not only precise but also rapid and efficient.

However, BMIs activated by MTL neurons may not be practical. Apart from technical and safety concerns with long-term implants, these neurons are not ideal for the task since the representation by them may be dynamic and of considerably plasticity. Thus, building a BMI based on the ability of the patient to think of a specific concept will prove futile once the neuron modifies its preference. BMI research might be more efficient in regions such as motor brain areas, which may have more stable representation based on population and not sparse coding. That said, the results from single unit studies in the MTL and frontal regions can provide significantly improved methods for testing BMIs with human subjects and for the improvement of the techniques used to read out the neural signal for future BMIs based on nonmotor regions (Naci et al., 2013).

Reading Thoughts and Urges during Their Formation

Recording the activity of individual nerve cells and identifying the pattern of activity underlying a particular percept can result not only in the real-time decoding of present thoughts but also in the prediction of coming thoughts. As the results of the free-recall study (Gelbard-Sagiv et al., 2008) suggest, it may be possible to trace neuronal activity that occurs before the actual experienced or reported recall (see figure 11.2).

In a study inspired by the classic Libet experiment from the 1980s (Libet et al., 1983), Fried and colleagues were able to isolate a set of neurons whose firing reliably predicted volition (Fried et al., 2011). Fried asked patients to look at a one-dial clock, which was revolving rapidly (roughly 2.2 s for a revolution). Subjects were asked to let the clock complete one full revolution before choosing, at their own will, to press a key (notated "P" for "Press"). Then subjects were asked to point to the time where they "felt ~~th~~urge" to perform the movement. Thus, subjects were asked to separate the moment of the actual action from the perceived moment of the decision or will to act (notated "W" for "Will"). Subjects timed this moment at around 200 ms before the actual button press.

Fried and colleagues were able to identify individual neurons, mostly in the supplementary and presupplementary motor areas and in anterior cingulate cortex, which became active up to 1000 ms prior to this estimated urge moment. The results from this study not only show that

one can predict an action before it was consciously recognized but also show the mechanisms of such volition. The main finding suggests an increase in the number of neurons involved in the encoding of the movement or the intention to move as the urge moment approaches. Of the 760 medial frontal neurons recorded in 12 subjects, 55 changed their activity 1000 ms before the urge moment, and the number increased to 128 neurons in the 400-ms window before this moment. A linear decoder based on the population activity could detect significant changes in activity 500 ms before the urge moment on 90% of the trials and 1000 ms before on 70% of the trials (see figure 11.6).

The ability to decode an action before its realization on a single trial basis, as seen here, is unique. For comparison, most works in fMRI or EEG, including that of Libet, required the averaging of dozens of trials in order to achieve similar level of accuracy in predicting the will. These results not only shed light on one of the most intriguing questions in neuroscience and philosophy, pertaining to the notion of free will and its underlying neuronal mechanisms, but also demonstrate access to information that is not necessarily available to the subject.

The Future: Accessing Conscious and Unconscious Thoughts

What makes one pattern of neuronal activity a part of our own subjective experience while other patterns remain “locked” in our brain and never reach our conscious awareness? Single spike resolution, together with the ease of getting reports on subjective experience, makes neurophysiological recording in humans a unique and singular setup to explore this type of question.

Being able to both ask subjects about their conscious thoughts, as well as independently read their formation, directly from the brain, has fundamental implications for understanding the processes that lead to the formation of conscious percepts and thoughts. Additionally, this direct readout of brain activity can access those percepts and thoughts which are processed by the brain but remain inaccessible and unconscious. Here, methods like binocular rivalry, flash suppression, change blindness, or attentional manipulations that separate conscious perception from external stimulations can prove extremely helpful.

Similarly, another interesting aspect of the process of human recollection that one can investigate is its associative nature. Thoughts and recollections are triggered many times by related or associated stimuli, which are later used to generate predictions (Bar, 2007). The network of associations is being built throughout our life based on our experiences and is unique to each person. Working with humans who can easily provide us with an aperture to their own network of associations enables us to see this process unfold in terms of neuronal activity. As we are able to see an individual concept being reflected upon, and occasionally even identify an association related to it, we can move on to study the formation of an association or the learning of a new concept. After we create an association between concepts in the patient’s mind, we can study the speed of the interaction, the timing and latency of the neuronal responses to each associated concept, and the similarity between multiple subjects in this associations buildup.

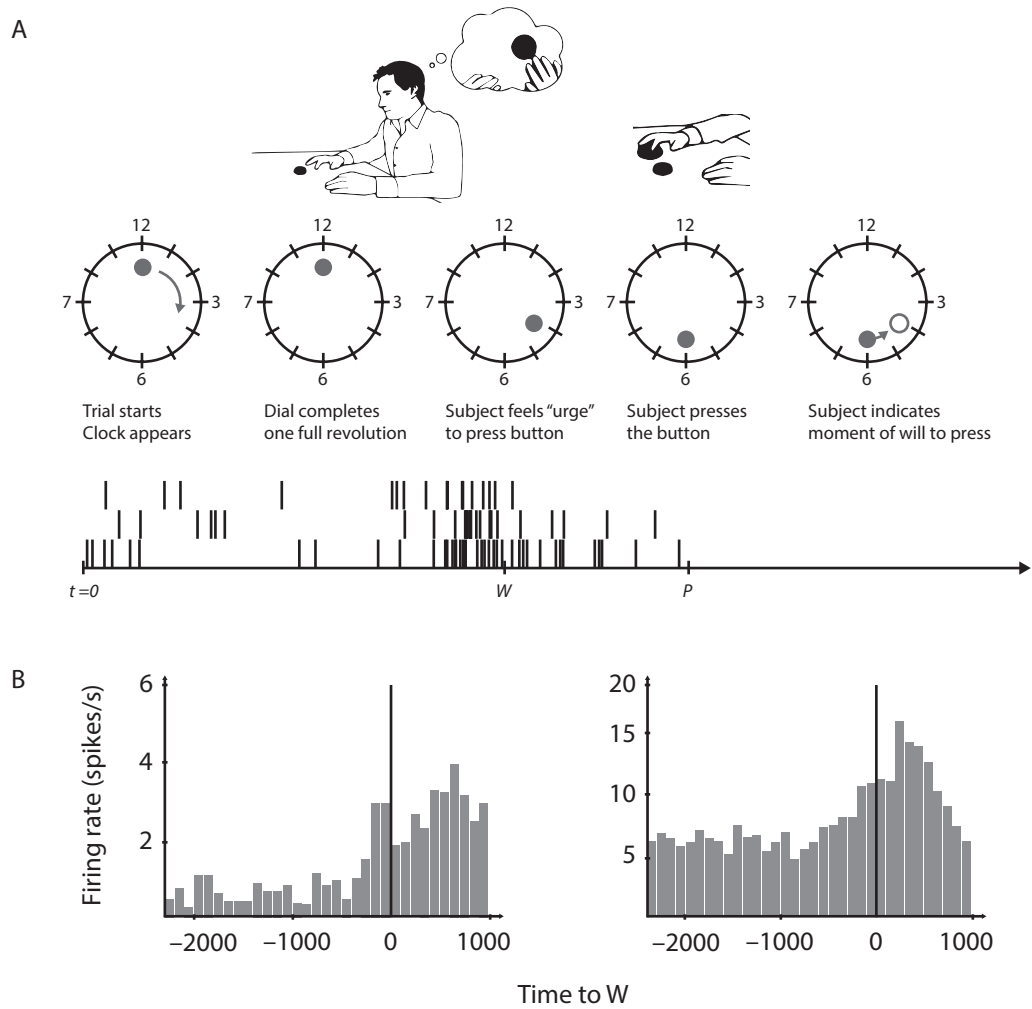


Figure 11.6
 (A) Schematic diagram depicting the experimental paradigm (Libet et al., 1983). Subjects were shown an analog clock and were asked to press a key with their right index finger, at will, anytime after one rotation of the clock. After the key press event (P), the clock dial stopped and subjects were asked to indicate the time of onset of the "urge/decision" to press the key (W). (B) Histograms showing the responses of a neuron (left) in the left anterior cingulate displaying a significant response after W (rank sum testp < 10^{-6}) and a neuron (right) in the left presupplementary motor area with response onset prior to W (rank sum testp < 10^{-3}). All plots are aligned to W (time = 0).

Finally, the ability to read thoughts from patterns of neuronal activity may open a window to the reading of the thoughts that run through our mind when we are asleep, our dreams. Ultimately, when we study dreams, we are interested in accessing the internal states that generate the dreams—having access to the single neuron level during sleep can serve as a unique platform to conduct such studies. One way to study this will involve an earlier “learning phase” during wakefulness, in which an algorithm learns to decode percepts and imagery of a large number of concepts. Then the algorithm will use this database to try to pick up these same activity patterns during sleep. If the database is extensive enough, we can reconstruct the narrative of the dream from the recorded activity and compare it with patients’ reports when they wake up. A recent study (Shibata et al., 2013) has demonstrated such mechanisms for dream decoding using fMRI in humans. The spatiotemporal nature of fMRI limited this study to the ability to decode high-level categories of dream, such as the general and gist content of the dream. The group was able to tell whether the subject was seeing a person, a landmark, or objects in a given moment but was not able to increase the resolution to the level of the specificity of the person, the landmark, or the object. They surmise that the main limitation was merely the resolution of the imaging device. If this is true, then with the high resolution provided by single neuron recordings we might be able to provide a more detailed interpretation of dream content and begin to understand how are dreams similar and how they differ from real-life experiences.

In summary, we are now in an opportune and exciting time where understanding of the way thoughts are constructed in our brain, alongside the high spatiotemporal resolution of single neuron recordings, allows further exploration of some of the most intriguing questions of neuroscience pertaining to those very faculties which make us human.

References

- Bar, M. (2007). The proactive brain: Using analogies and associations to generate predictions. *Trends in Cognitive Sciences* 11, 280–289.
- Cerf, M., Thiruvengadam, N., Mormann, F., Kraskov, A., Quiroz Quiroga, R., Koch, C., et al. (2010). On-line, voluntary control of human temporal lobe neurons. *Nature*, 467, 1104–1108.
- Freud, S. (1954). The interpretation of dreams. In J. Strachey (Ed. & Trans.), *The standard edition of the complete psychological works of Sigmund Freud* (one vol. reprint of Vols. 4–5). London: Hogarth Press. (Original work published 1900)
- Fried, I., Mukamel, R., & Kreiman, G. (2011). Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. *Neuron*, 69, 548–562.
- Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., & Fried, I. (2008). Internally generated reactivation of single neurons in human hippocampus during free recall. *Science*, 322, 96–101.
- Haynes, J.-D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews. Neuroscience* 7, 523–534.
- Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., et al. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485, 372–375.
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience* 8, 679–685.
- Kay, K. N., Naselaris, T., Prenger, R. J., & Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, 452, 352–355.

- Kreiman, G., Koch, C., & Fried, I. (2000). Imagery neurons in the human brain. *Nature*, 408, 357–361.
- Lebedev, M. A., & Nicolelis, M. A. L. (2006). Brain-machine interfaces: Past, present and future. *Trends in Neurosciences*, 29, 536–546.
- Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential): The unconscious initiation of a freely voluntary act. *Brain*, 106, 623–642.
- Naci, L., Cusack, R., Jia, V. Z., & Owen, A. M. (2013). The brain's silent messenger: Using selective attention to decode human thought for brain-based communication. *Journal of Neuroscience*, 33, 9385–9393.
- Nishimoto, S., Vu, A. T., Naselaris, T., Benjamini, Y., Yu, B., & Gallant, J. L. (2011). Reconstructing visual experiences from brain activity evoked by natural movies. *Current Biology*, 21, 1641–1646.
- Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Sciences*, 10, 424–430.
- Nyberg, L., Habib, R., McIntosh, A. R., & Tulving, E. (2000). Reactivation of encoding-related brain activity during memory retrieval. *Proceedings of the National Academy of Sciences of the United States of America*, 97, 11120–11124.
- Polyn, S. M., Natu, V. S., Cohen, J. D., & Norman, K. A. (2005). Category-specific cortical activity precedes retrieval during memory search. *Science*, 310, 1963–1966.
- Quiroga, R. Q., Reddy, L., Koch, C., & Fried, I. (2007). Decoding visual inputs from multiple neurons in the human temporal lobe. *Journal of Neurophysiology*, 98, 1997–2007.
- Rees, G., Kreiman, G., & Koch, C. (2002). Neural correlates of consciousness in humans. *Nature Reviews. Neuroscience*, 3, 261–270.
- Shibata, K., Watanabe, T., Sasaki, Y., & Kawato, M. (2011). Perceptual learning incepted by decoded fMRI neurofeedback without stimulus presentation. *Science*, 334, 1413–1415.
- Velliste, M., Perel, S., Spalding, M. C., Whitford, A. S., & Schwartz, A. B. (2008). Cortical control of a prosthetic arm for self-feeding. *Nature*, 453, 1098–1101.
- Wheeler, M. E., Petersen, S. E., & Buckner, R. L. (2000). Memory's echo: Vivid remembering reactivates sensory-specific cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 97, 11125–11129.