

Progressions

Uncovering a Timescale Hierarchy by Studying the Brain in a Natural Context

Uri Hasson

Psychology Department and the Neuroscience Institute, Princeton University, Princeton, New Jersey 08544

As we train multiple generations of students to narrowly design clever, carefully controlled experiments in our confined lab spaces, we may fail to notice, as a field, that we have overlooked fundamental aspects of human cognition. This is a first-person account of how our research and understanding of the neural code were forever transformed when we decided to open the lab's door to the natural world. This journey started with the decision to shift from controlled stimuli to natural dynamic and “messy” stimuli. This transition enabled us to focus on how information is accumulated and processed over time. As a result, we have discovered a new topographic mapping of the hierarchy of cortical processing timescales. I will conclude with a general observation of the paradigm shift occurring in the field as it increasingly emphasizes the study of the neural processes that underlie human behavior in natural, everyday contexts. I am excited to share this journey with you.

“The present contains nothing more than the past, and what is found in the effect was already in the cause,”
Henri Bergson, *Creative Evolution*.

The Shift from the Lab to the Natural World

My journey started in 1999 when I trained as a PhD student in Rafi Malach's lab at the Weizmann Institute. Rafi's lab was vibrant and creative. These were the early days of the cognitive neuroscience revolution, and the new imaging tools let us map the previously unknown landscape of high-level visual processing areas (Levy et al., 2001; Hasson et al., 2002). Nancy Kanwisher had just discovered the fusiform face area (Kanwisher et al., 1997) and the parahippocampal place area (Epstein et al., 1999). However, as our research progressed and more and more of our papers were accepted, I became increasingly concerned about the ecological validity of our findings.

In each of our studies, we were zealous about controlling as many variables as possible while carefully manipulating our dimensions of interest. For example, we minimized the effect of eye movements by having our subjects fixate on a small red point. By normalizing the images, we controlled for luminance- and contrast-induced variations in neural activity. We removed the impact of colors by applying grayscale filtering. We often resorted to using line drawings to remove the effect of texture.

Received Dec. 17, 2024; revised Jan. 3, 2025; accepted Jan. 7, 2025.

I would like to express my heartfelt gratitude to all the extraordinary individuals with whom I have had the privilege to collaborate over the years. First and foremost, I want to thank my PhD advisor, Professor Rafi Malach, whose guidance and mentorship have been instrumental in shaping me into the scientist I am today. I am also profoundly grateful to my postdoctoral advisors, the late Professor Nava Rubin and Professor David Heeger, who, together with me, discovered the temporal processing hierarchy. A special thanks go to my fantastic lab members and colleagues at Princeton University, with whom I worked on further charting the processing temporal hierarchy: Dr. Yulia Lerner, Dr. Chris Honey, Dr. Janice Chen, Dr. Claire Chang, Dr. Mor Regev, Dr. Greg Stephens, Dr. Yarra Yeshurun, and Professor Kenneth A. Norman. I would also like to thank Dr. Chris Honey and Dr. Liat Hasenfratz for their insightful comments on this paper.

The authors declare no competing financial interests.

Correspondence should be addressed to Uri Hasson at hasson@princeton.edu.

<https://doi.org/10.1523/JNEUROSCI.2368-24.2025>

Copyright © 2025 the authors

We presented isolated objects and faces cropped on a gray background to remove the effect of spatial context. Finally, we continually presented the objects and faces, one at a time, using event-related designs that removed temporal dynamics and memory-induced interactions across images. We hoped back then, and perhaps we still hope to this day, that gradually, by manipulating a few variables at a time, using an incremental divide-and-conquer strategy, we could collectively, one day, aggregate all of these piecemeal studies into a coherent and rich neurocomputational model of the human brain. But what if the brain responds in entirely different, even opposing, ways in the natural world, where myriad variables interact in time in complex and nonlinear ways?

Ecological Validity and the Replications Crisis

Ecological validity poses significant challenges in our field, exacerbating and compounding the replication crisis (Pashler and Harris, 2012; Shamay-Tsoory and Mendelsohn, 2019). To replicate a study, we are asked to carefully follow all the steps and recipes implemented in the original studies. Any tiny deviation from the original setup may be used to explain our failure to replicate a study. However, the more significant issue here is that the request to replicate a study verbatim fails to address the issue of ecological validity. Even upon replication, what is the significance of an effect so frail as to only materialize in a confined and artificial set of parameters (not representing the real world)?

Robustness and generalization across natural contexts provide a new framework for assessing the ecological validity of our findings. In real life, we do not have the luxury of eliminating confounds. Instead, the brain's task is to detect, select, and amplify the relevant dimensions as a function of context while all parameters naturally vary. Therefore, a system designed for real-life object recognition should handle ever-changing brightness, contrast, color, and context fluctuations. For example, one tool the visual system uses to filter variability is eye movements, which help scan, attend to, and process regions of interest in the natural scenes. Therefore, instead of removing or

fixing dimensions, our experimental goal should be to show that our findings hold across all these contexts.

In the early 2000s, we were still uncertain whether the FFA and PPA would maintain functional selectivity as subjects viewed dynamic natural images freely. As a result, I decided to extend my PhD for a year to undertake a final project to test the selectivity of the cortical (including visual) systems during the processing of complex audio-visual movies. This decision changed my career forever.

Stepping Out into the Natural World and Uncovering Intersubject Correlation

At that time, the idea of imaging the brain as it processes real-life situations was unheard of. Since we were all trained to control every aspect of our studies, it was challenging for us to relinquish control. We spent days arguing about allowing subjects to watch a movie freely, concerned about the lack of control over eye movements. Similarly, we debated whether to include or mute the soundtrack, as we were concerned that the interaction between different sensory modalities would complicate the analyses and interpretations. We also discussed the dialogues, plot, and the potential for the brain to adapt to recurring elements as the movie unfolded. Ultimately, we decided to take the risks and scan people as they freely watched a 30 min audiovisual segment of *The Good, the Bad and the Ugly* by Sergio Leone (Hasson et al., 2004).

Furthermore, back then most fMRI studies relied on general linear models and event-related designs, considered cutting-edge analyses. However, these methods were unsuitable for analyzing the continuous movie-watching dataset. Thus, we needed to develop new ways to analyze our data.

As we relinquished control, we anticipated that each individual would have idiosyncratic ways of viewing and interpreting the movie, resulting in significant uncontrolled and unaccounted variations in neural activity among individuals. However, to my surprise, as I sampled voxels from all subjects across the visual system, including V1, V2, V4, LO, FFA, and PPA, a very different pattern started to emerge. The response patterns in each of these visual regions were highly selective and distinct from the responses in other visual areas. Within each area, however, the responses across subjects were highly similar (correlated), indicating that variability across subjects was low (and not high as a priori expected) as they watched the movie. Excited about these results, I rushed into Rafi's office. We decided to run a comprehensive whole-brain analysis in search of all voxels that responded similarly across subjects as they watched the movie. To do so, we measured the correlation of time courses across all subjects in each voxel across the entire cortex, which we later named intersubject correlation (ISC, Nastase et al., 2019).

The ISC measures the reliability of activity over time across subjects rather than measuring the average activity level as tested by event-related analysis (Hasson et al., 2010). To our astonishment, the selective alignment across subjects was widespread (Hasson et al., 2004). We thus observed selective yet shared correlated activity across subjects in visual areas, auditory areas along the superior temporal gyrus, language areas including Wernicke's and Broca's areas, and many high-level parietal and frontal areas associated with the default mode network (DMN).

Functional Selectivity in the Natural World

We were curious about what drove the selective yet shared activity across subjects in each brain region as they watched the

movie. We wondered if our results aligned with the expected patterns of selectivity observed in previous controlled lab-based experiments or if we had uncovered new organizational patterns. Our investigation showed evidence of both. Most importantly, we discovered a new topographical mapping of processing timescale that had yet to be identified.

To our relief, the FFA, PPA, and LO retained their known spatial selectivity to faces, places, and objects, respectively, as predicted by the more controlled experiments (Kanwisher, 2010). Our naturalistic approach showed that the functional selectivity of high-order visual areas was robust to variations in contrast, luminance, perspective, and motion. Moreover, when a scene is composed of multiple objects—for example, when Clint Eastwood rides a horse around a deserted town while carrying a gun in search of the bad guys—we found that the selectivity was determined by attention and eye movements. Recall our extensive discussions about the risk of losing control over subjects' gaze as we contemplated removing the fixation. In hindsight, in taking the risk, we made a wise decision, as the film effectively guided all our subjects' gazes, naturally directing them to focus on the same parts of the scene. Using this complex natural stimulus and allowing for spontaneous behavior, we uncovered the shared alignment in eye movements and the parallel and synchronous enhancement in neural selectivity while disparate individuals watched the movie.

After our initial study with *The Good, the Bad and the Ugly*, we expanded our research to test the generalizability (i.e., ecological validity) of our findings using a variety of natural stimuli. We found that a strong correlation among subjects can also occur in everyday situations despite not resembling carefully crafted Hollywood films with multimillion-dollar budgets. For instance, sharing personal accounts during storytelling events could induce high ISCs in language and high-order areas across listeners (Nastase et al., 2021). Additionally, we found that high ISCs can emerge when ordinary people recall their daily memories (Chen et al., 2017) and engage in everyday conversations (Goldstein et al., 2025) and that these effects are modulated in people with certain neuropsychiatric disorders (Kronberg et al., 2024). Finally, we discovered a coupling between the brain responses of a speaker and a listener as they engaged in natural conversations (Stephens et al., 2010; Silbert et al., 2014). As a result, we determined that alignment across subjects captures shared neural signals associated with how social groups process and share information (Hasson et al., 2012; Zada et al., 2024).

Temporal Context Constantly Shapes Our Minds Over Time

As we delved into the complexity and multidimensional nature of our movie and storytelling stimuli and the corresponding neural activity, we soon realized that transitioning to natural contexts raises new questions that are easy to overlook in lab settings. In a typical event-related design, researchers present isolated context-less stimuli, one at a time, for a brief duration. For example, in many of our studies, we presented isolated images of faces, objects, and houses, one at a time, for 500 milliseconds each (Hasson et al., 2001; Malach et al., 2002). Then, we calculated the average neural response for images of a particular category (e.g., faces) compared with the average neural activity for the other categories of stimuli. Any variability in the neural responses among stimuli within the same category was generally considered a nuisance to average out and minimize; time was not modeled.

However, the neural variability over time among stimuli is essential rather than a nuisance, as it captures the contextual aspect of natural stimuli. To illustrate, a 30 min movie comprises ~43,200 unique images projected at 24 frames per second. These images are meticulously edited and fused with a soundtrack of environmental sounds, dialogues, and music designed to tell a cohesive story. In such rich contexts, the meaning of any frame or sound bite at a specific time in the movie can be shaped by complex, nonlinear interactions with all preceding and subsequent audiovisual stimuli. Similarly, in a spoken story, each word derives its full meaning from the subtle contextual interactions with all other words in the narrative. Such dynamic, ever-changing, temporal interactions, the core of natural stimuli, are deliberately removed from traditional event-related experiments conducted in a lab setting.

Trigger-averaging analyses cannot capture the context-dependent richness imbued in natural dynamic stimuli (Ben Yaakov et al., 2012). For example, averaging images with a close-up of faces in a movie removes any neural signal associated with our ability to recognize different faces in natural contexts robustly. After all, the neural response to Clint Eastwood (the “good guy”) should not be washed out by averaging it with the neural response to Eli Wallach (the “ugly guy”). Furthermore, subtle nested temporal dependencies would be lost if one is to characterize the “typical” neural response, even for a single face in a movie or a single word in a narrative. In each scene, Clint Eastwood’s face conveys a subtly different expression that the brain picks up. Averaging all close-ups of Clint’s face will wash away such signals. Similarly, each word can change meaning in each scene based on context. Take the word “cold,” for instance: the phrase “you are cold as ice” could refer to your body temperature or your perceived personality trait, depending on the broader temporal context of the conversation. Thus, averaging all occurrences of a word or a word category (such as nouns versus verbs) will likely eliminate subtle context-dependent interactions. These context-dependent interactions compelled us to rethink how the brain processes and integrates dynamic, context-sensitive information across multiple timescales.

The Processing Timescale Hierarchy

During my postdoctoral research at NYU, I collaborated with Professors Nava Rubin and David Heeger to explore how past events shape the processing of incoming information across the cortical hierarchy. Our approach involved manipulating the temporal structure of movies and audiobooks at various temporal granularities without oversimplifying or reducing the dimensionality of the stimuli or compromising the ecological validity of our research studies.

Similar to spatial receptive fields, which enable neurons and visual areas to integrate information across space, we posited that neurons and brain regions possess temporal integration windows that allow them to integrate information over time. To characterize the processing time scale for each cortical area, we measured how traces of prior events (recent memory) influenced moment-to-moment neural activity (online processes) during minutes-long real-life stimuli. To do so, we manipulated the temporal structure of movies across multiple timescales. We began by dividing each stimulus into smaller temporal units, segmenting each movie into individual frames, short clips of a few seconds between cuts, and 30–40 s continuous plot fragments (Hasson et al., 2008). In a follow-up experiment at Princeton, where I started my lab, we applied a similar approach to a spoken story. We segmented the story into individual sound bites and

broke it down further into words, sentences, and paragraphs (Lerner et al., 2011). Next, the temporal structure of the stimulus was varied by scrambling the order of the units at each temporal granularity while maintaining it within each time unit. We also played the movies and audiobooks in reverse order to measure neural activity at the single keyframe (shorter timescale) level. In the backward condition, the past context in the forward intact condition became the future context while maintaining the temporal integrity of each keyframe.

Like the known spatial receptive field topography, the temporal integration window increases from early sensory areas, which have a millisecond integration window, to high-order areas, which have a minutes-long integration window. For example, a clear pattern was observed in the visual areas as we moved along the cortical hierarchy. In early visual areas, the temporal integration windows were only a few milliseconds long. In contrast, higher-order visual areas had integration windows lasting several seconds, while high-level processing regions in the frontal and parietal cortices had integration windows extending to tens of seconds (Hasson et al., 2008). A similar hierarchical pattern was noted in the auditory, language, and high-order areas when we scrambled the audiobooks (Fig. 1A; Lerner et al., 2011).

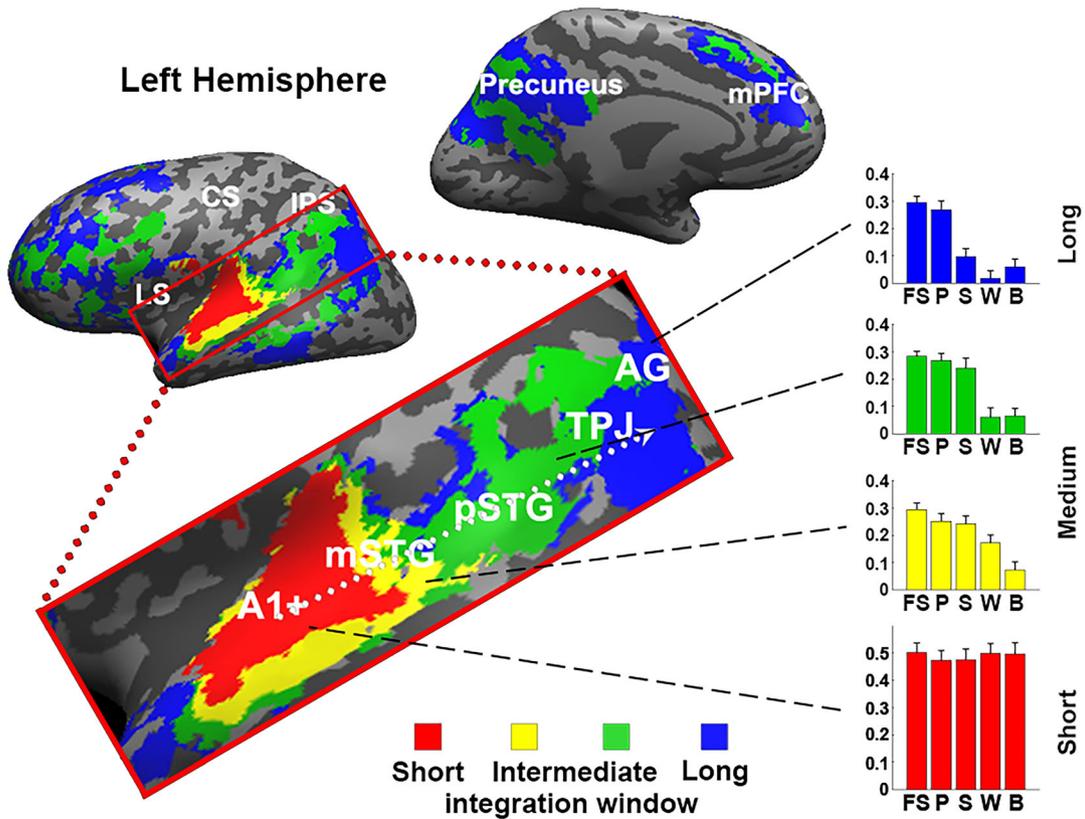
The Struggle to Publish rather than Perish

It was challenging to publish our first paper on the hierarchy of processing timescales along the visual system. While many prominent researchers were excited about exploring neural dynamics during natural-world processing, others heavily criticized the work. The idea of relinquishing control and using varied multidimensional stimuli contradicted some core beliefs about how to conduct rigorous research. I remember David remarking that our work was causing division in the field and that he had never seen such emotionally charged reviews before. After more than 10 rounds of revisions, during which we incorporated additional control experiments and analyses to support our findings, I almost gave up. I worried that my career might hit a dead-end. Fortunately, the editors at *JNeurosci* stepped in and were willing to bring in new reviewers to re-evaluate the previous reviews and guide the process until the paper was published (Hasson et al., 2008). To our relief, the *JNeurosci* readily accepted the second paper, which describes the temporal integration window along the auditory-to-language hierarchy (Fig. 1A; Lerner et al., 2011).

Characterization of the Cortical Processing Hierarchy

Over the years, my lab at Princeton has further characterized the cortical processing hierarchy with the help of my exceptionally talented students and postdocs. We examined how information flows along the processing timescale hierarchy during natural communication (Chang et al., 2022; Goldstein et al., in press), music perception (Farbood et al., 2015; Piazza et al., 2021), and narrative comprehension (Yeshurun et al., 2017). Honey et al. (2012) reported a timescale hierarchy using intracranial recordings (Fig. 1B) and showed how the persistence of neural dynamics (Chang et al., 2021) could be measured using autocorrelation and spectral methods (Honey et al., 2012; Stephens et al., 2013). Additionally, we investigated the interactions between high-level cortical areas with minute-long temporal integration windows and the hippocampus (Chen et al., 2015; Zuo et al., 2020). We also explored the role of attention in modulating these processes (Regev et al., 2019). Lastly, we studied how information is

A) Process memory timescales measured using fMRI



B) Process memory timescales measured using ECoG

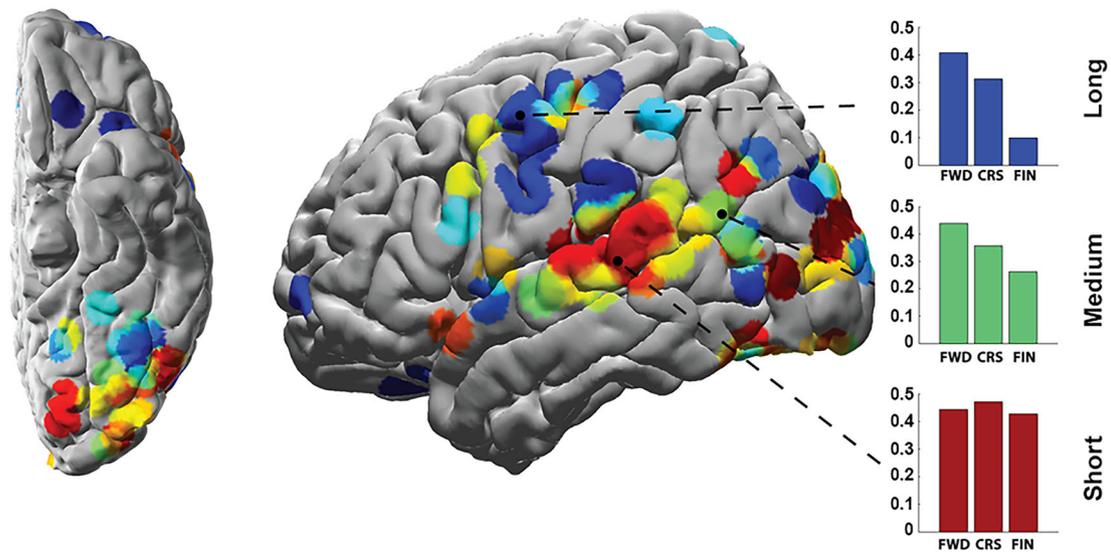


Figure 1. Hierarchical of temporal processing timescales. **A**, fMRI map of the gradual transition from short to long temporal integration windows along the temporal–parietal axis mapped using audio narratives. The color of each voxel indicates the shortest timescale of coherence in the stimulus that produced a reliable intersubject response (red, story played backward; yellow, story with word-order scrambled; green, story with sentence-order scrambled; blue, story with paragraph-order scrambled). fMRI time courses in early auditory areas (A1+) were reliable across subjects exposed to the same stimulus; this was true at all scrambling levels, from the intact full story (FS) to scrambled paragraphs (P), scrambled sentences (S), scrambled words (W), and backward speech (B). Further up the processing hierarchy, more and more stimulus history affected responses in the present moment. At the top of the hierarchy, areas such as the temporal parietal junction (TPJ) responded reliably only at the full story and paragraph levels. **B**, Electroencephalography (ECoG) map of the gradual transition from short to long temporal integration windows mapped using an audiovisual movie. Shorter integration windows were predominantly found near primary sensory areas, while longer integration windows were found further away from sensory areas. Early auditory areas (A1+) responded reliably across all scrambling levels, from the intact full movie (FWD) to the coarse scrambled movie (CRS) and the scrambled movie (FIN). Further up the processing hierarchy, more and more stimulus history affected responses in the present moment. At the top of the hierarchy, areas such as the lateral prefrontal cortex responded with much greater reliability at the intact and coarse scrambled movie levels. Figure adapted from Hasson et al. (2015).

organized into events and scaled along the processing hierarchy (Lerner et al., 2014; Baldassano et al., 2017).

This discovery of the processing timescale topography opened a new line of research. Many labs further explored the processing timescale hierarchy using numerous techniques, including fMRI, EcOG, EEG, and single-unit measurements. The processing hierarchy was also replicated across species, from humans to primates (Chaudhuri et al., 2015) to rodents (Rudelt et al., 2024). Murray et al. (2014) revealed a timescale hierarchy even in single-neuron recordings. The notion of a timescale hierarchy was integrated into the broader mapping in the primate cerebral cortex: the gradient from external to internal cortical networks (Margulies et al., 2016); gradients of gene expression, myelination, and cortical thickness (Burt et al., 2018; Shafiei et al., 2020); as well as gradients of brain network connectedness (Baria et al., 2013).

Moreover, the timescale gradient was extended to a whole brain perspective, including subcortical circuits (Raut et al., 2020). It was also reported to persist across task states in behaving monkeys (Manea et al., 2024) and humans (Wolff et al., 2022) and became essential in many theories of learning (Bernacchia et al., 2011; Soltani et al., 2021). Going forward, an important question is whether the timescale hierarchy is an inherent property of brain circuits—essentially a genetically encoded

foundation that shapes and influences learning—or if it emerges from the process of learning the multi-scale statistics of the natural world (Hasson et al., 2020). Several studies have investigated the effectiveness of using a timescale hierarchy for sequence learning models (Chung et al., 2016) and encoding models of neural dynamics (Vo et al., 2023).

Processing Timescale and Memory Systems

The temporal integration window measures how past events impact the processing of incoming information in each cortical region. For the past to affect the processing of incoming information, a form of memory trace must linger and shape the processing at each level of the temporal processing hierarchy. However, we were perplexed about how to connect the concept of the temporal integration window with established memory constructs. Is a temporal integration window associated with short-term memory, working memory, or long-term memory? This question puzzled my postdoc, Christopher Honey, and me for weeks. It is difficult to reconcile the idea that each brain area can accumulate information over a given temporal window with the notion that we possess distinct memory systems—separate from processing systems—that store information over short and long timescales. After many weeks of internal discussions, we finally realized that the idea of a temporal integration window compels us to unite the

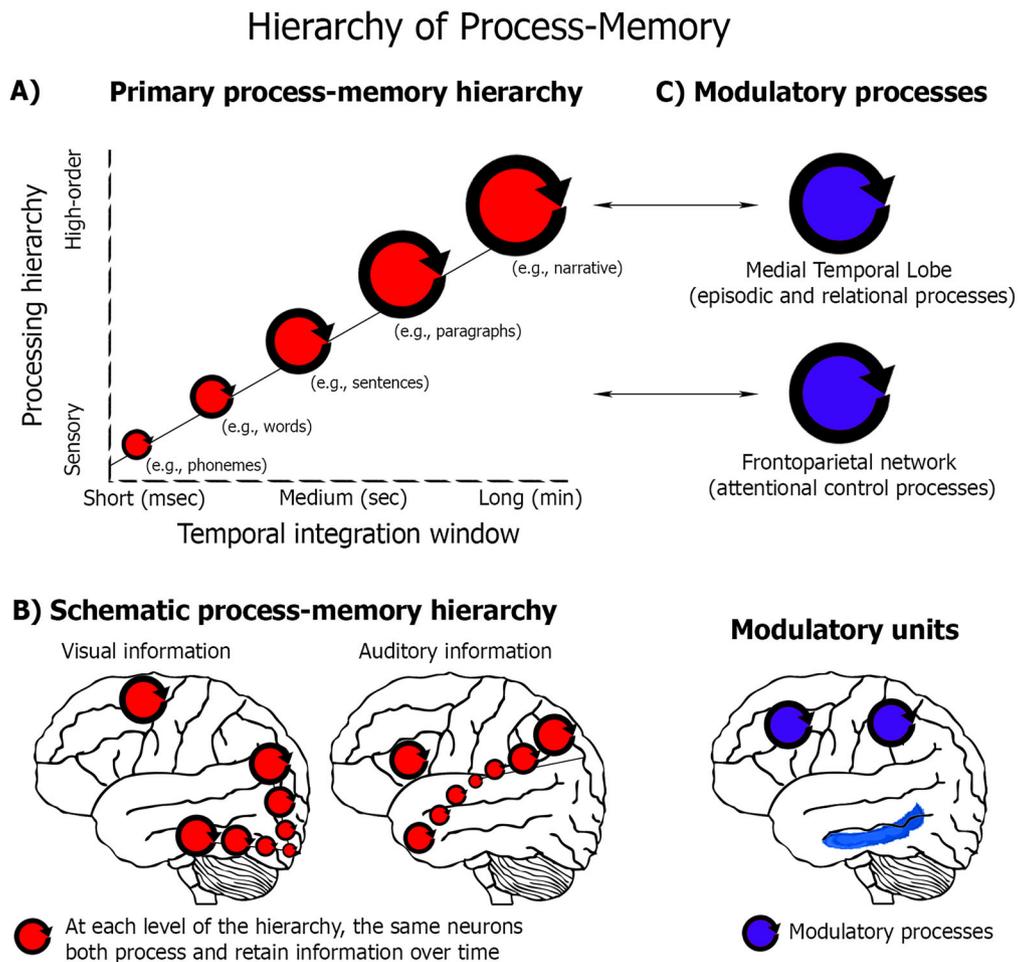


Figure 2. *A*, A hierarchy of process memory framework. Memory is integral to the operation of each cortical area, and there is no separation between the processing and information storage units. Furthermore, each region’s processing timescale (operationalized by measuring the temporal integration window) increases in a topographically organized manner, from milliseconds in early sensory areas to minutes in high-order areas. *B*, A schematic process memory hierarchy for auditory and visual stimulation (for actual data, Fig. 1). *C*, Primary versus modulatory process memory. Two additional processes (blue circles) modulate the primary process memories hierarchy (red circles): attentional control and episodic memory processes. Figure adapted from Hasson et al. (2015).

concepts of working memory, long-term memory, and neural processing into one cohesive framework rather than dissociate them. In essence, the processing timescale views memory as a fundamental part of any information processing circuit. After all, each neuron functions as both a processing unit, dynamically synthesizing information from its dendrites, and a memory unit, capable of modifying its synaptic connections with other neurons. Similarly, each neural circuit accumulates (memorizes) and synthesizes (processes) information over its preferred timescale. Thus, temporal respective windows integrate the accumulation of memories over time together with online information processing in a unified process-memory framework (Hasson et al., 2015).

Specifically, while all cortical circuits have the process-memory capacity to store information over time, the temporal integration window increases hierarchically from early sensory areas to higher-order perceptual and cognitive regions (Fig. 2). Early sensory areas have a short temporal integration window (tens of milliseconds), which allows them to integrate sensory information (e.g., a few phonemes to recognize a word). Mid-level regions (such as language areas) have a medium memory integration window (several seconds), enabling the integration of information into longer sequences (e.g., integrating words while analyzing a sentence). At the top of the processing hierarchy, the DMN areas have a long memory integration window (seconds to minutes) required to integrate information (e.g., sentences) over long episodes (e.g., as we engage in long conversations, audiobooks, or movies). Finally, the DMN has strong connections with the medial temporal lobe and hippocampus, allowing it to increase the contextual window by adding episodic information accumulated over minutes, days, and even years (Chen et al., 2016; Yeshurun et al., 2021).

The recent success of artificial neural networks in processing natural stimuli, such as visual, auditory, and language stimuli, sheds additional light on our process-memory hypothesis. Like biological neural networks, artificial neural networks dismiss the classical notion of segregation between processing units and memory units. Specifically, in artificial neural networks, each neuron functions as a processing unit and a memory unit that can update its connectivity weights. Furthermore, instead of using a short-term memory cache, as in Von Neumann's digital computers, artificial neural networks can use recurrent activity or attention heads to retain contextual information from previous events. Finally, we showed that neural activity in cortical areas at the top of the processing hierarchy can be modulated by information gathered over a very long period: by augmenting large language models with additional episodic storage, such as memory-augmented large language models, relevant prior contexts (encoded and retrieved over days to years) change processing of a given event (Wang et al., 2023). The recent success in building deep models that can process the natural world offers a new computational framework for understanding the neural code that underlies the processing of the natural world in the human brain (Tikochinski et al., 2025).

Concluding Remarks

Using natural stimuli, we initially set out to test the ecological validity and reliability of our laboratory-based findings and theories in real-life situations. However, we soon realized that controlled experiments cannot accurately capture the brain's neural dynamics when processing real-world experiences. We conclude our journey by considering ecological validity as a fundamental guiding principle for developing cognitive theories

and testing and interpreting our hypotheses. Findings that can only be replicated under narrow, artificial, and highly controlled conditions, which do not materialize in the real world, may lack significance and relevance to everyday cognition, and should be questioned. When we started using movies and audiobooks in our studies, the use of natural stimuli to explore the neural basis of everyday cognitive processes was at its inception. Over the years, we have observed how natural stimuli have led to a paradigm shift in how we theorize and conduct our experiments in neuroscience. Finally, with the recent advancements in deep generative models that can process natural stimuli and interact with the real world without simplifying or controlling any input aspect, natural neuroscience is on the brink of a second paradigm shift. This shift is triggering a new set of model-based studies that may lead us to develop better theories and computational models for understanding how the human brain develops and operates in the real world (Hasson et al., 2020; Nastase et al., 2020).

References

- Baldassano C, Chen J, Zadbood A, Pillow JW, Hasson U, Norman KA (2017) Discovering event structure in continuous narrative perception and memory. *Neuron* 95:709–721.e5.
- Baria AT, Mansour A, Huang L, Baliki MN, Cecchi GA, Mesulam MM, Apkarian AV (2013) Linking human brain local activity fluctuations to structural and functional network architectures. *Neuroimage* 73:144–155.
- Ben-Yakov A, Honey CJ, Lerner Y, Hasson U (2012) Loss of reliable temporal structure in event-related averaging of naturalistic stimuli. *Neuroimage* 63:501–506.
- Bernacchia A, Seo H, Lee D, Wang X-J (2011) A reservoir of time constants for memory traces in cortical neurons. *Nat Neurosci* 14:366–372.
- Burt JB, Demirtaş M, Eckner WJ, Navejar NM, Ji JL, Martin WJ, Bernacchia A, Anticevic A, Murray JD (2018) Hierarchy of transcriptomic specialization across human cortex captured by structural neuroimaging topography. *Nat Neurosci* 21:1251–1259.
- Chang CHC, Lazaridi C, Yeshurun Y, Norman KA, Hasson U (2021) Relating the past with the present: information integration and segregation during ongoing narrative processing. *J Cogn Neurosci* 33:1106–1128.
- Chang CHC, Nastase SA, Hasson U (2022) Information flow across the cortical timescale hierarchy during narrative construction. *Proc Natl Acad Sci U S A* 119:e2209307119.
- Chaudhuri R, Knoblauch K, Gariel M-A, Kennedy H, Wang X-J (2015) A large-scale circuit mechanism for hierarchical dynamical processing in the primate cortex. *Neuron* 88:419–431.
- Chen J, Hasson U, Honey CJ (2015) Processing timescales as an organizing principle for primate cortex. *Neuron* 88:244–246.
- Chen J, Honey CJ, Simony E, Arcaro MJ, Norman KA, Hasson U (2016) Accessing real-life episodic information from minutes versus hours earlier modulates hippocampal and high-order cortical dynamics. *Cereb Cortex* 26:3428–3441.
- Chen J, Leong YC, Honey CJ, Yong CH, Norman KA, Hasson U (2017) Shared memories reveal shared structure in neural activity across individuals. *Nat Neurosci* 20:115–125.
- Chung J, Ahn S, Bengio Y (2016) Hierarchical multiscale recurrent neural networks. In *arXiv [cs.LG]*. arXiv: <https://openreview.net/pdf?id=S1di0sfgl>
- Epstein R, Harris A, Stanley D, Kanwisher N (1999) The parahippocampal place area: recognition, navigation, or encoding? *Neuron* 23:115–125.
- Farbood MM, Heeger DJ, Marcus G, Hasson U, Lerner Y (2015) The neural processing of hierarchical structure in music and speech at different timescales. *Front Neurosci* 9:157.
- Goldstein A, et al. (2025) A unified acoustic-to-speech-to-language embedding space captures the neural basis of natural language processing in everyday conversations. *Nat Hum Behav* [in press].
- Hasson U, Chen J, Honey CJ (2015) Hierarchical process memory: memory as an integral component of information processing. *Trends Cogn Sci* 19:304–313.
- Hasson U, Ghazanfar AA, Galantucci B, Garrod S, Keysers C (2012) Brain-to-brain coupling: a mechanism for creating and sharing a social world. *Trends Cogn Sci* 16:114–121.

- Hasson U, Hendler T, Bashat D, Malach R (2001) Vase or face? A neural correlate of shape-selective grouping processes in the human brain. *Neuroimage* 13:744–753.
- Hasson U, Levy I, Behrmann M, Hendler T, Malach R (2002) Eccentricity bias as an organizing principle for human high-order object areas. *Neuron* 34:479–490.
- Hasson U, Malach R, Heeger DJ (2010) Reliability of cortical activity during natural stimulation. *Trends Cogn Sci* 14:40–48.
- Hasson U, Nastase SA, Goldstein A (2020) Direct fit to nature: an evolutionary perspective on biological and artificial neural networks. *Neuron* 105:416–434.
- Hasson U, Nir Y, Levy I, Fuhrmann G, Malach R (2004) Intersubject synchronization of cortical activity during natural vision. *Science* 303:1634–1640.
- Hasson U, Yang E, Vallines I, Heeger D, Rubin N (2008) A hierarchy of temporal receptive windows in human cortex. *J Neurosci* 28:2539–2550.
- Honey CJ, Thesen T, Donner TH, Silbert LJ, Carlson CE, Devinsky O, Doyle WK, Rubin N, Heeger DJ, Hasson U (2012) Slow cortical dynamics and the accumulation of information over long timescales. *Neuron* 76:423–434.
- Kanwisher N (2010) Functional specificity in the human brain: a window into the functional architecture of the mind. *Proc Natl Acad Sci U S A* 107:11163–11170.
- Kanwisher N, McDermott J, Chun MM (1997) The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J Neurosci* 17:4302–4311.
- Kronberg G, Ceceli AO, Huang Y, Gaudreault P-O, King SG, McClain N, Alia-Klein N, Goldstein RZ (2024) Shared orbitofrontal dynamics to a drug-themed movie track craving and recovery in heroin addiction. *Brain:awae369*.
- Lerner Y, Honey CJ, Katkov M, Hasson U (2014) Temporal scaling of neural responses to compressed and dilated natural speech. *J Neurophysiol* 111:2433–2444.
- Lerner Y, Honey C, Silbert LJ, Hasson U (2011) Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *J Neurosci* 31:2906–2915.
- Levy I, Hasson U, Avidan G, Hendler T, Malach R (2001) Center-periphery organization of human object areas. *Nat Neurosci* 4:533–539.
- Malach R, Levy I, Hasson U (2002) The topography of high-order human object areas. *Trends Cogn Sci* 6:176–184.
- Manea AMG, Maisson DJ-N, Voloh B, Zilverstand A, Hayden B, Zimmermann J (2024) Neural timescales reflect behavioral demands in freely moving rhesus macaques. *Nat Commun* 15:2151.
- Margulies DS, et al. (2016) Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proc Natl Acad Sci U S A* 113:12574–12579.
- Murray JD, et al. (2014) A hierarchy of intrinsic timescales across primate cortex. *Nat Neurosci* 17:1661–1663.
- Nastase SA, et al. (2021) The “narratives” fMRI dataset for evaluating models of naturalistic language comprehension. *Sci Data* 8:250.
- Nastase SA, Gazzola V, Hasson U, Keysers C (2019) Measuring shared responses across subjects using intersubject correlation. *Soc Cogn Affect Neurosci* 14:667–685.
- Nastase SA, Goldstein A, Hasson U (2020) Keep it real: rethinking the primacy of experimental control in cognitive neuroscience. *Neuroimage* 222:117254.
- Pashler H, Harris CR (2012) Is the replicability crisis overblown? Three arguments examined. *Perspect Psychol Sci* 7:531–536.
- Piazza EA, Nencheva ML, Lew-Williams C (2021) The development of communication across timescales. *Curr Dir Psychol Sci* 30:459–467.
- Raut RV, Snyder AZ, Raichle ME (2020) Hierarchical dynamics as a macroscopic organizing principle of the human brain. *Proc Natl Acad Sci U S A* 117:20890–20897.
- Regev M, Simony E, Lee K, Tan KM, Chen J, Hasson U (2019) Propagation of information along the cortical hierarchy as a function of attention while reading and listening to stories. *Cereb Cortex* 29:4017–4034.
- Rudelt L, González Marx D, Spitzner FP, Cramer B, Zierenberg J, Priesemann V (2024) Signatures of hierarchical temporal processing in the mouse visual system. *PLoS Comput Biol* 20:e1012355.
- Shafiei G, Markello RD, Vos de Wael R, Bernhardt BC, Fulcher BD, Misis B (2020) Topographic gradients of intrinsic dynamics across neocortex. *Elife* 9:e62116.
- Shamay-Tsoory SG, Mendelsohn A (2019) Real-life neuroscience: an ecological approach to brain and behavior research. *Perspect Psychol Sci* 14:841–859.
- Silbert LJ, Honey CJ, Simony E, Poeppel D, Hasson U (2014) Coupled neural systems underlie the production and comprehension of naturalistic narrative speech. *Proc Natl Acad Sci U S A* 111:E4687–96.
- Soltani A, Murray JD, Seo H, Lee D (2021) Timescales of cognition in the brain. *Curr Opin Behav Sci* 41:30–37.
- Stephens GJ, Honey CJ, Hasson U (2013) A place for time: the spatiotemporal structure of neural dynamics during natural audition. *J Neurophysiol* 110:2019–2026.
- Stephens GJ, Silbert LJ, Hasson U (2010) Speaker-listener neural coupling underlies successful communication. *Proc Natl Acad Sci U S A* 107:14425–14430.
- Tikochinski R, Goldstein A, Meiri Y, Hasson U, Reichart R (2025) Incremental accumulation of linguistic context in artificial and biological neural networks. *Nat Commun* 16:803.
- Vo VA, Jain S, Beckage N, Chien HYS, Obinwa C, Huth AG (2023) A unifying computational account of temporal context effects in language across the human cortex. In *bioRxiv*.
- Wang W, Dong L, Cheng H, Liu X, Yan X, Gao J, Wei F (2023) Augmenting language models with long-term memory. In: *Advances in neural information processing systems* (Oh A, Naumann T, Globerson A, Saenko K, Hardt M, Levine S, eds), pp 74530–74543. Curran Associates, Inc.
- Wolff A, Berberian N, Golesorkhi M, Gomez-Pilar J, Zilio F, Northoff G (2022) Intrinsic neural timescales: temporal integration and segregation. *Trends Cogn Sci* 26:159–173.
- Yeshurun Y, Nguyen M, Hasson U (2017) Amplification of local changes along the timescale processing hierarchy. *Proc Natl Acad Sci U S A* 114:9475–9480.
- Yeshurun Y, Nguyen M, Hasson U (2021) The default mode network: where the idiosyncratic self meets the shared social world. *Nat Rev Neurosci* 22:181–192.
- Zada Z, et al. (2024) A shared model-based linguistic space for transmitting our thoughts from brain to brain in natural conversations. *Neuron* 112:3211–3222.e5.
- Zuo X, Honey CJ, Barense MD, Crombie D, Norman KA, Hasson U, Chen J (2020) Temporal integration of narrative information in a hippocampal amnesic patient. *Neuroimage* 213:116658.