Postgraduate Certificate in AI and Cognitive Psychology

Cognitive Psychology Principles

Attention – the selective allocation of cognitive resources to specific information while ignoring others.
Related terms: selective attention, sustained attention, divided attention, attentional bias.
Explanation: Attention functions as a spotlight, enhancing processing of attended stimuli and suppressing unattended inputs. In cognitive psychology, models such as the filter theory and capacity theory describe how limited processing resources are distributed.
Example: When a driver focuses on traffic signals, peripheral objects like billboards receive less processing.
Practical applications: Designing user interfaces that minimize extraneous visual clutter to reduce attentional overload; training AI vision systems to prioritize salient features using attention mechanisms.
Challenges: Measuring covert attention objectively; accounting for individual differences in attentional capacity; integrating top-down goals with bottom-up stimulus salience in computational models.

Associative Learning – a process whereby a relationship between two stimuli or a stimulus and a response is formed.
Related terms: classical conditioning, operant conditioning, contingency, reinforcement schedule.
Explanation: Through repeated pairings, the organism learns that one event predicts another, leading to anticipatory responses. Pavlov's dogs salivating to a bell illustrates classical conditioning, while Skinner's lever-pressing demonstrates operant conditioning.
Example: A user learns that clicking a "download" button results in a file transfer, reinforcing the behavior.
Practical applications: Human-computer interaction (HCI) designs that employ reward feedback to shape user behavior; AI agents that use reinforcement learning algorithms mirroring associative principles.
Challenges: Distinguishing between correlation and causation in complex environments; dealing with extinction when reinforcement is removed; modeling the timing of stimulus presentation accurately.

Automaticity – the ability to perform tasks with little conscious effort after extensive practice.
Related terms: procedural memory, skill acquisition, chunking, cognitive load.
Explanation: As a skill becomes automatic, processing shifts from controlled, effortful mechanisms to fast, unconscious pathways, freeing working-memory capacity for other tasks.
Example: Experienced typists can type without looking at the keyboard, allowing them to focus on content.
Practical applications: Designing training programs that aim for automaticity to improve efficiency; implementing AI modules that pre-process routine tasks to reduce computational load.
Challenges: Detecting when a skill has reached automaticity; preventing over-automation that leads to loss of situational awareness; managing interference when dual tasks compete for residual resources.

Bayesian Brain – the hypothesis that the brain interprets sensory information by combining prior knowledge with incoming data in a probabilistic manner.

Related terms: predictive coding, prior probability, posterior distribution, hierarchical inference.
Explanation: According to this view, perception is an inferential process where expectations (priors) are updated with sensory evidence to generate the most probable interpretation (posterior).
Example: In a noisy environment, listeners use linguistic context to fill in missing phonemes, effectively "guessing" the intended words.
Practical applications: Developing AI perception systems that incorporate Bayesian updating for robust decision-making under uncertainty; improving adaptive user interfaces that predict user intent.
Challenges: Quantifying priors in real-world settings; computational complexity of hierarchical Bayesian models; reconciling the theory with neurophysiological data.

Cognitive Architecture – a theoretical framework that specifies the structural components and processes underlying human cognition.
Related terms: ACT-R, SOAR, production system, mental modules.
Explanation: Cognitive architectures model how perception, memory, reasoning, and action are coordinated, often using symbolic representations and rule-based processing. They provide a bridge between psychological theory and computational implementation.
Example: ACT-R simulates human multitasking by modeling the interaction of declarative memory and procedural modules.
Practical applications: Building cognitive agents that emulate human problem-solving strategies; evaluating educational interventions by simulating learner behavior.
Challenges: Balancing model fidelity with computational tractability; integrating subsymbolic processes such as neural network learning; validating architectures against diverse empirical data.

Cognitive Bias – systematic patterns of deviation from rational judgment that arise from mental shortcuts.
Related terms: confirmation bias, anchoring effect, availability heuristic, framing effect.
Explanation: Biases emerge because the brain economizes on effort, leading to predictable errors in perception, memory, and decision-making. They influence how information is encoded, retrieved, and evaluated.
Example: Investors may overweight recent market trends (availability) when forecasting future performance.
Practical applications: Designing AI decision-support tools that flag potential human biases; creating debiasing training for clinicians to improve diagnostic accuracy.
Challenges: Identifying subtle biases in complex tasks; distinguishing bias from legitimate heuristics; mitigating bias without impairing efficient processing.

Cognitive Load – the total amount of mental effort being used in working memory at any given time.
Related terms: intrinsic load, extraneous load, germane load, capacity theory.
Explanation: Cognitive load theory posits that instructional design should manage the limited working-memory capacity to optimize learning. Intrinsic load reflects task complexity, extraneous load stems from poor design, and germane load supports schema construction.
Example: A tutorial that presents a single algorithm step at a time reduces extraneous load compared with a

dense slide showing the entire code.
Practical applications: Developing e-learning modules that segment information; informing AI tutoring systems to adapt content based on real-time load estimates.
Challenges: Measuring cognitive load objectively (e.g., via pupillometry); accounting for individual differences in capacity; maintaining engagement while minimizing overload.

Cognitive Map – an internal representation of spatial relationships in the environment.
Related terms: mental navigation, place cells, grid cells, spatial memory.
Explanation: Cognitive maps enable flexible navigation, allowing individuals to infer novel routes and remember landmarks. In rodents, place cells fire when the animal occupies a specific location, supporting map formation.
Example: A tourist uses a mental map of a city to find a new restaurant without a GPS.
Practical applications: Designing autonomous robots that construct and use cognitive maps for path planning; developing VR training that enhances spatial memory.
Challenges: Translating neural findings into computational models; handling dynamic environments where landmarks change; integrating multimodal cues (visual, auditory) into unified maps.

Cognitive Neuroscience – the interdisciplinary study of how brain structures and functions give rise to mental processes.
Related terms: fMRI, EEG, lesion studies, neural correlates.
Explanation: By linking behavioral data with neuroimaging and electrophysiological measures, cognitive neuroscience seeks to map cognitive functions onto specific neural circuits.
Example: Functional MRI reveals activation of the prefrontal cortex during working-memory tasks.
Practical applications: Informing the design of brain-inspired AI architectures; guiding neurofeedback interventions for attention disorders.
Challenges: Inferring causality from correlational data; resolving the spatial-temporal trade-off in measurement techniques; integrating findings across scales from synapses to systems.

Cognitive Load Theory – a framework describing how instructional design influences the mental effort required for learning.
Related terms: schema acquisition, split-attention effect, modality effect, expertise reversal.
Explanation: The theory emphasizes that working memory has limited capacity; effective instruction reduces extraneous load and promotes germane processing to facilitate long-term memory formation.
Example: Pairing narration with visuals (modality effect) can lower load compared with text-only explanations.
Practical applications: Adaptive e-learning platforms that monitor learner load and adjust presentation style; AI-driven tutoring that sequences content based on mastery.
Challenges: Differentiating between intrinsic and extraneous load in complex domains; developing reliable real-time load metrics; preventing “expertise reversal” where novice-oriented strategies hinder advanced learners.

Dual-Process Theory – a model proposing two distinct systems for reasoning: a fast, automatic System 1 and a slower, deliberative System 2.
Related terms: heuristic processing, analytic reasoning, cognitive control, metacognition.
Explanation: System 1 operates effortlessly using heuristics and associative links; System 2 engages when tasks demand logical analysis, often overriding intuitive responses.
Example: Solving a simple arithmetic problem (2 + 2) taps System 1, whereas solving a complex probability puzzle requires System 2.
Practical applications: Designing AI agents that switch between rapid pattern recognition and deeper symbolic reasoning; creating educational tools that train learners to engage System 2 when appropriate.
Challenges: Determining the neural substrates of each system; preventing overreliance on System 1 leading to bias; modeling the interaction dynamics computationally.

Embodied Cognition – the view that cognitive processes are deeply rooted in the body's interactions with the physical world.
Related terms: sensorimotor grounding, affordances, situated cognition, body schema.
Explanation: According to this perspective, perception, action, and thought are inseparable; cognition emerges from real-time sensorimotor experiences rather than abstract symbol manipulation alone.
Example: Manipulating a puzzle piece physically helps a person understand spatial relationships better than visualizing it alone.
Practical applications: Developing robotics that learn through embodied exploration; enhancing virtual reality training by coupling perception with motor actions.
Challenges: Translating embodied principles into purely computational models; accounting for cultural and linguistic influences that modulate embodiment; measuring the contribution of bodily states to abstract reasoning.

Episodic Memory – the ability to recall personal experiences situated in time and place.
Related terms: autobiographical memory, hippocampus, retrieval cue, episodic future thinking.
Explanation: Episodic memories are encoded with contextual details, allowing mental time travel to relive past events or imagine future scenarios. The hippocampus plays a crucial role in binding these details.
Example: Remembering a birthday party includes the cake's flavor, the music, and the people present.
Practical applications: Designing AI assistants that can retrieve user-specific past interactions to personalize recommendations; therapeutic interventions for memory disorders that use cue-based retrieval.
Challenges: Differentiating episodic from semantic memory in neuroimaging; preventing false memories during reconstruction; modeling the vividness and subjective confidence of recollections.

Executive Functions – higher-order cognitive processes that regulate goal-directed behavior.
Related terms: inhibition, cognitive flexibility, planning, dorsolateral prefrontal cortex.
Explanation: Executive functions encompass the ability to suppress inappropriate responses, switch between tasks, and maintain future goals, orchestrating other mental operations.
Example: While driving, a driver inhibits the impulse to check a phone, thereby maintaining focus on the

road.
Practical applications: Implementing AI systems that monitor user workload and adapt task demands; creating cognitive training programs to strengthen executive control in aging populations.
Challenges: Isolating distinct executive components in behavioral tests; assessing transfer effects of training to real-world tasks; integrating executive function models with low-level perceptual processes.

Feature Integration Theory – a model of visual attention proposing that basic visual features are first processed in parallel, then combined into coherent objects through focused attention.
Related terms: preattentive processing, binding problem, visual search, conjunction search.
Explanation: The theory distinguishes a “feature-search” stage where attributes like color or orientation are detected automatically, and a “conjunction-search” stage requiring attention to bind features into object representations.
Example: Finding a red circle among green circles is rapid (feature search), whereas locating a red circle among red squares and green circles is slower (conjunction search).
Practical applications: Designing visual dashboards that exploit preattentive features for quick status detection; informing computer vision algorithms about hierarchical attention mechanisms.
Challenges: Explaining instances where conjunction searches are performed faster than predicted; accounting for top-down influences on feature selection; modeling the neural basis of binding.

Gestalt Principles – a set of laws describing how the visual system organizes elements into groups or unified wholes.
Related terms: proximity, similarity, continuity, closure, figure-ground.
Explanation: The brain spontaneously groups stimuli based on relational properties, facilitating efficient perception of patterns and objects.
Example: A series of aligned dots is perceived as a line due to the principle of continuity.
Practical applications: User-interface design that leverages proximity to indicate related controls; AI image segmentation that incorporates Gestalt constraints for more human-like grouping.
Challenges: Quantifying the relative strength of competing principles; extending principles to multimodal (audio-visual) integration; integrating Gestalt insights with deep-learning feature hierarchies.

Heuristics – mental shortcuts that simplify decision-making but can lead to systematic errors.
Related terms: availability heuristic, representativeness heuristic, satisficing, fast-and-frugal heuristics.
Explanation: Heuristics reduce cognitive load by allowing rapid judgments based on limited information, often relying on salient cues or prior experience.
Example: Estimating the probability of a plane crash by recalling recent news stories (availability).
Practical applications: Implementing heuristic-based AI agents for real-time resource allocation; training professionals to recognize when heuristic use may be inappropriate.
Challenges: Balancing speed and accuracy; identifying contexts where heuristics are beneficial versus detrimental; formalizing heuristics within computational models.

Implicit Memory – unconscious memory systems that influence behavior without conscious recollection.

Related terms: procedural memory, priming, skill learning, basal ganglia.
Explanation: Implicit memory manifests as improved performance on tasks through prior exposure, even when the individual cannot explicitly describe the prior experience.
Example: Typing a familiar word faster after having seen it earlier, despite not remembering the exposure.
Practical applications: Designing adaptive interfaces that subtly prime users toward desired actions; leveraging implicit learning in rehabilitation for motor disorders.
Challenges: Measuring implicit memory separate from explicit recall; preventing unintended priming effects that bias user decisions; modeling implicit processes in AI without explicit representations.

Information Processing Model – a framework depicting cognition as a sequence of stages: encoding, storage, and retrieval.
Related terms: sensory memory, short-term memory, long-term memory, serial processing.
Explanation: The model likens the mind to a computer, where information passes through buffers, is transformed, and later accessed, providing a basis for experimental paradigms.
Example: Visual information first enters iconic memory, then moves to working memory for manipulation, and finally consolidates into long-term storage.
Practical applications: Structuring instructional design to align with encoding and retrieval phases; building AI pipelines that mimic staged processing for natural-language understanding.
Challenges: Accounting for parallel, interactive processes observed in the brain; integrating affective and motivational factors; reconciling the model with evidence of distributed neural representations.

Metacognition – awareness and regulation of one’s own cognitive processes.
Related terms: self-monitoring, metacognitive judgments, reflective practice, learning strategies.
Explanation: Metacognition involves evaluating the adequacy of knowledge, planning how to approach tasks, and adjusting strategies based on feedback. It plays a pivotal role in effective learning and problem solving.
Example: A student assesses confidence after answering a question and decides whether to review the material.
Practical applications: AI tutoring systems that prompt learners to reflect on their confidence levels; workplace training that incorporates metacognitive prompts to improve decision quality.
Challenges: Accurately assessing metacognitive accuracy; distinguishing between overconfidence and appropriate self-assessment; embedding metacognitive loops in autonomous agents.

Neural Plasticity – the brain’s capacity to reorganize its structure and function in response to experience.
Related terms: long-term potentiation, synaptic pruning, experience-dependent plasticity, neurogenesis.
Explanation: Plastic changes occur at synaptic, circuit, and systems levels, enabling learning, memory formation, and recovery after injury.
Example: Musicians develop enlarged cortical representations of finger movements compared with non-musicians.
Practical applications: Designing brain-computer interfaces that adapt to user neural changes; informing AI

systems that continuously update network weights in response to environmental feedback.
Challenges: Mapping specific behavioral changes to underlying neural mechanisms; balancing stability and flexibility in models of plasticity; translating animal findings to human cognition.

Neural Networks (Biological) – interconnected groups of neurons that process information through collective dynamics.
Related terms: excitatory/inhibitory balance, receptive fields, spike-timing dependent plasticity, cortical columns.
Explanation: Biological neural networks exhibit parallel processing, distributed representations, and adaptive learning, inspiring artificial neural network architectures.
Example: The visual cortex contains orientation-selective neurons that respond preferentially to edges of specific angles.
Practical applications: Developing neuromorphic hardware that mimics cortical processing efficiency; using insights from cortical connectivity to improve deep-learning architectures.
Challenges: Capturing the stochastic nature of spiking activity in deterministic models; scaling biologically realistic simulations; bridging the gap between low-level neuronal dynamics and high-level cognition.

Object Recognition – the ability to identify and categorize visual entities despite variations in viewpoint, lighting, and occlusion.
Related terms: invariant representation, feature hierarchy, template matching, ventral stream.
Explanation: Recognition relies on extracting stable features (edges, textures) and integrating them into abstract representations that support categorization across conditions.
Example: Recognizing a car whether it is seen from the front or side.
Practical applications: Autonomous vehicles that must detect pedestrians and traffic signs; AI systems that assist visually impaired users by labeling objects in real time.
Challenges: Achieving robustness to novel variations; reducing reliance on massive labeled datasets; explaining how recognition decisions are formed (interpretability).

Procedural Memory – a type of implicit memory for skills and actions that can be performed without conscious awareness.
Related terms: motor learning, skill acquisition, basal ganglia, habit formation.
Explanation: Procedural memory encodes sequences of movements and procedural knowledge, allowing fluent execution after extensive practice.
Example: Riding a bicycle involves balance and coordination that are hard to verbalize.
Practical applications: Rehabilitation programs that exploit procedural learning for motor recovery; AI agents that acquire motor skills through reinforcement learning mirroring human habit formation.
Challenges: Measuring procedural memory independently of explicit knowledge; addressing interference between competing motor sequences; modeling the transition from effortful learning to automatic execution.

Psychophysical Scaling – methods for quantifying the relationship between physical stimulus intensity and

perceived magnitude.
Related terms: Weber’s law, Fechner’s law, just-noticeable difference (JND), magnitude estimation.
Explanation: Scaling establishes how changes in stimulus parameters (e.g., brightness) translate to subjective experience, often revealing logarithmic relationships.
Example: Doubling the intensity of a sound does not double perceived loudness; a larger increase is required for the same perceptual effect.
Practical applications: Calibrating audio-visual displays to align with human perception; informing AI that generates multimedia content to match user sensitivity thresholds.
Challenges: Accounting for individual variability; extending scaling to complex, multidimensional stimuli; integrating scaling laws into predictive models of perception.

Reinforcement Learning (RL) – a computational framework where agents learn to maximize cumulative reward through trial-and-error interactions with an environment.
Related terms: reward prediction error, Q-learning, policy gradient, exploration-exploitation.
Explanation: RL mirrors associative learning concepts, using value functions to estimate future returns and updating them based on observed outcomes, analogous to dopamine-driven learning in the brain.
Example: A robot learns to navigate a maze by receiving positive feedback when reaching the goal and negative feedback for collisions.
Practical applications: Training autonomous systems for complex decision-making; modeling human learning in tasks such as game playing or skill acquisition.
Challenges: Balancing exploration of new strategies with exploitation of known good actions; dealing with sparse or delayed rewards; ensuring stability and convergence in high-dimensional spaces.

Schema – organized knowledge structures that represent generic concepts and guide information processing.
Related terms: script, mental model, top-down processing, assimilation.
Explanation: Schemas provide expectations about typical features and relationships, influencing perception, memory encoding, and retrieval. They enable efficient inference but can also lead to distortions.
Example: The “restaurant” schema includes expectations of being seated, ordering food, and paying the bill.
Practical applications: Designing chatbots that activate appropriate conversational scripts based on user context; educational tools that help learners build accurate schemas in complex domains.
Challenges: Updating schemas when confronted with contradictory information; measuring the influence of schemas on memory accuracy; preventing overgeneralization that hinders flexible reasoning.

Semantic Memory – the long-term storage of factual knowledge about the world, independent of personal experience.
Related terms: concepts, categories, knowledge network, anterior temporal lobe.
Explanation: Semantic memory organizes information into interconnected webs, supporting language comprehension, reasoning, and problem solving.
Example: Knowing that Paris is the capital of France is a semantic fact.

Practical applications: Knowledge-base construction for AI that requires factual retrieval; designing curricula that reinforce semantic networks for deeper understanding.
Challenges: Differentiating semantic from episodic traces in neuroimaging; dealing with semantic decay or distortion over time; modeling the flexible retrieval of related concepts.

Signal Detection Theory (SDT) – a statistical approach to measuring the ability to discriminate signal from noise.
Related terms: hits, false alarms, d′ (d-prime), criterion, receiver operating characteristic (ROC).
Explanation: SDT separates sensitivity (ability to detect a signal) from decision bias (tendency to respond "yes" or "no"), providing a nuanced view of performance.
Example: In a medical screening test, SDT quantifies how well clinicians detect disease (hits) while accounting for false positives.
Practical applications: Evaluating AI classifiers for balanced accuracy; training users to adjust decision criteria based on risk contexts.
Challenges: Estimating parameters reliably with limited data; extending SDT to multidimensional decision spaces; integrating subjective confidence measures.

Social Cognition – the study of how people process, store, and apply information about others and social situations.
Related terms: theory of mind, empathy, attribution, stereotype, mirror neurons.
Explanation: Social cognition involves interpreting intentions, emotions, and behaviors, often using specialized neural circuits that support perspective-taking and affective resonance.
Example: Inferring that a colleague is frustrated based on facial expression and tone.
Practical applications: Designing socially aware AI companions that can recognize and respond to human emotions; training programs that improve interpersonal skills by enhancing theory-of-mind abilities.
Challenges: Modeling nuanced, context-dependent social judgments; avoiding bias in AI that mirrors human stereotypes; measuring internal states like empathy objectively.

Symbolic Reasoning – a computational approach that manipulates explicit symbols and rules to derive conclusions.
Related terms: production rules, logic programming, knowledge representation, deductive inference.
Explanation: Symbolic systems encode knowledge in discrete symbols and apply formal operations, enabling transparent, interpretable reasoning processes.
Example: An expert system uses if-then rules to diagnose equipment failures.
Practical applications: Building AI that can explain its decisions through logical proof steps; integrating symbolic modules with neural components for hybrid reasoning.
Challenges: Capturing commonsense knowledge that is often vague; scaling rule bases without combinatorial explosion; interfacing symbolic representations with sub-symbolic perception.

Theory of Mind (ToM) – the capacity to attribute mental states—beliefs, desires, intentions—to oneself and others.

Related terms: false-belief task, mentalizing network, perspective-taking, empathy.
Explanation: ToM enables prediction of behavior by modeling how others perceive and intend, relying on brain regions such as the temporoparietal junction and medial prefrontal cortex.
Example: Understanding that a friend who cannot see the hidden card will make a different choice than you.
Practical applications: Developing conversational agents that anticipate user needs; creating virtual training environments that assess ToM development in children.
Challenges: Measuring ToM in non-verbal populations; ensuring AI agents do not infer incorrect mental states; integrating ToM reasoning with real-time perception.

Working Memory – a limited-capacity system for temporarily holding and manipulating information.
Related terms: phonological loop, visuospatial sketchpad, central executive, capacity limits.
Explanation: Working memory supports complex cognition such as reasoning, comprehension, and learning by maintaining active representations while other processes operate.
Example: Solving a mental arithmetic problem requires holding intermediate results in working memory.
Practical applications: Designing adaptive interfaces that offload information to reduce working-memory load; creating AI agents that simulate human-like short-term storage for dialogue management.
Challenges: Determining the exact capacity (often cited as 7 ± 2 items) across modalities; modeling the dynamic interaction between the central executive and subsidiary stores; accounting for individual differences in working-memory span.

Visual Attention – the selective processing of visual information, directing gaze and neural resources toward relevant stimuli.
Related terms: saccades, fixation, covert attention, attentional spotlight.
Explanation: Visual attention involves both overt movements (eye-movements) and covert shifts (attention without eye-movement), guided by saliency, task goals, and expectations.
Example: While reading, the eyes fixate on each word, but attention may also anticipate the next word based on context.
Practical applications: Eye-tracking interfaces that adapt content based on fixation patterns; AI vision systems that prioritize regions of interest to improve efficiency.
Challenges: Disentangling bottom-up salience from top-down goal influences; modeling the temporal dynamics of attentional shifts; handling attentional capture by sudden, irrelevant stimuli.

Zone of Proximal Development (ZPD) – the range of tasks that a learner can perform with guidance but not yet independently.
Related terms: scaffolding, Vygotskian theory, assisted learning, developmental readiness.
Explanation: The ZPD concept emphasizes the importance of social interaction and support in advancing cognitive abilities, highlighting the role of more knowledgeable others.
Example: A child can solve a puzzle with hints from a teacher but cannot complete it alone.
Practical applications: Adaptive tutoring systems that provide incremental hints within a learner's ZPD; designing collaborative learning environments that pair novices with experts.

Challenges: Accurately assessing the boundaries of a learner's ZPD; avoiding over-scaffolding that hampers autonomous skill development; translating the concept into automated personalization algorithms.

Neurofeedback – a technique that provides real-time information about neural activity, enabling individuals to self-regulate brain states.
Related terms: EEG biofeedback, operant conditioning, brain-computer interface, self-regulation.
Explanation: By presenting feedback (e.g., visual bar representing alpha power), users learn to modulate specific neural patterns, leading to improvements in attention, anxiety, or motor control.
Example: Training athletes to increase sensorimotor rhythm activity to enhance focus during competition.
Practical applications: Therapeutic interventions for ADHD that aim to normalize attentional networks; integrating neurofeedback into AI-driven wellness platforms.
Challenges: Ensuring specificity of training targets; addressing inter-individual variability in responsiveness; maintaining long-term transfer of learned regulation to everyday contexts.

Predictive Coding – a theoretical framework proposing that the brain constantly generates predictions about sensory input and updates them based on prediction errors.
Related terms: hierarchical inference, error minimization, generative models, free-energy principle.
Explanation: Higher-level areas send predictions downward, while lower-level areas compute mismatches (prediction errors) that propagate upward, refining internal models.
Example: Expecting a familiar song to continue in a certain key; when a note deviates, the brain registers a prediction error and updates its expectation.
Practical applications: Building AI systems that anticipate user actions, reducing latency in interaction; developing compression algorithms that encode data by predicting upcoming values.
Challenges: Quantifying prediction error signals in neural recordings; scaling hierarchical models to complex, real-world data; reconciling predictive coding with alternative accounts of perception.

Transfer Learning – the process of applying knowledge gained in one domain to improve learning or performance in another domain.
Related terms: domain adaptation, fine-tuning, knowledge reuse, cross-task generalization.
Explanation: In cognition, transfer occurs when skills or concepts learned in one context facilitate problem solving in a novel but related context.
Example: Learning to play piano can aid the acquisition of keyboard skills for computer programming.
Practical applications: Pre-training AI language models on large corpora and then fine-tuning for specific tasks; educational curricula that deliberately connect prior knowledge to new topics.
Challenges: Identifying conditions under which transfer is positive versus negative (interference); measuring the depth of transferred representations; preventing catastrophic forgetting when adapting to new domains.

Unconscious Processing – cognitive operations that occur without conscious awareness, influencing perception, decision-making, and behavior.
Related terms: subliminal perception, blindsight, implicit bias, automaticity.

Explanation: Unconscious processes can extract statistical regularities, guide motor actions, and shape preferences, often operating in parallel with conscious deliberation.
Example: A driver responds to a sudden brake light before consciously registering the stimulus.
Practical applications: Designing subtle cues in interfaces that nudge user behavior without overt instruction; employing unconscious priming to enhance learning retention.
Challenges: Detecting unconscious influences experimentally; distinguishing unconscious processing from low-level conscious awareness; ethical considerations of manipulating behavior covertly.

Visual Short-Term Memory (VSTM) – the temporary storage of visual information for a few seconds.
Related terms: iconic memory, capacity limit, change detection, spatial resolution.
Explanation: VSTM retains details of visual scenes, allowing comparison across brief intervals; its capacity is typically limited to 3-4 items with high fidelity.
Example: Remembering the location of a target among distractors in a visual search task.
Practical applications: Designing dashboards that present critical information within VSTM limits to avoid loss; developing AI models that simulate human-like short-term visual retention for rapid scene analysis.
Challenges: Measuring the precise capacity and decay rates; accounting for the influence of attention on VSTM performance; integrating VSTM models with longer-term memory systems.

Working Memory Training – systematic exercises aimed at enhancing the capacity and efficiency of working memory.
Related terms: n-back task, dual-task training, adaptive difficulty, cognitive enhancement.
Explanation: Training protocols adapt task difficulty based on performance, seeking neuroplastic changes that improve executive functions and fluid intelligence.
Example: An n-back app that increases the number of items to be remembered as the user improves.
Practical applications: Cognitive rehabilitation for individuals with traumatic brain injury; corporate programs to boost employee multitasking abilities.
Challenges: Demonstrating transfer of gains to untrained tasks; avoiding placebo effects; establishing long-term durability of improvements.

Symbol Grounding Problem – the challenge of connecting abstract symbols used in AI systems to meaningful real-world referents.
Related terms: embodiment, semantics, referential grounding, perceptual symbols.
Explanation: Without grounding, symbols remain meaningless; cognitive psychology suggests that sensorimotor experiences provide the basis for symbol meaning, informing AI approaches that link language to perception.
Example: The word “apple” acquires meaning through visual, gustatory, and tactile experiences.
Practical applications: Developing multimodal AI that learns word meanings from paired images and sounds; creating robots that associate commands with physical actions.
Challenges: Achieving robust grounding across varied contexts; preventing symbol drift as systems learn from noisy data; reconciling symbolic abstraction with continuous perceptual input.

Transfer of Training – the application of skills or knowledge learned in one setting to different tasks or environments.
Related terms: near transfer, far transfer, generalization, skill retention.
Explanation: Effective transfer depends on similarity of underlying principles, contextual cues, and the learner's ability to abstract core components.
Example: A pilot trained on a flight simulator applies procedural knowledge to a real aircraft.
Practical applications: Designing simulation-based curricula that maximize far transfer to real-world performance; evaluating AI agents for ability to generalize policies across domains.
Challenges: Predicting which aspects of training will generalize; mitigating negative transfer where prior learning interferes with new tasks; measuring transfer quantitatively.

Temporal Discounting – the tendency to devalue rewards that are delayed in time, favoring immediate gratification.
Related terms: delay aversion, hyperbolic discounting, impulsivity, intertemporal choice.
Explanation: Cognitive models describe discounting as a function that declines steeply for near-future rewards and more slowly for distant ones, reflecting a trade-off between present and future benefits.
Example: Choosing to receive $50 today over $100 in a month.
Practical applications: Designing incentive structures that encourage long-term adherence (e.g., health programs); incorporating discounting models into AI decision-making for realistic human behavior simulation.
Challenges: Capturing individual variability in discount rates; integrating affective states that modulate discounting; reconciling hyperbolic models with normative economic theories.

Self-Regulation – the ability to control thoughts, emotions, and behaviors to achieve goals.
Related terms: executive control, goal-directed behavior, emotion regulation, self-efficacy.
Explanation: Self-regulation draws on working memory, attentional control, and motivational processes to monitor progress and adjust actions.
Example: A student resists the urge to check social media while studying for an exam.
Practical applications: Building AI assistants that provide prompts to sustain focus; training interventions that improve self-control in addictive behaviors.
Challenges: Measuring self-regulation in ecologically valid settings; distinguishing between trait and state aspects; ensuring interventions do not induce excessive rigidity.

Neuroplasticity-Based AI – artificial intelligence models inspired by the brain's adaptive mechanisms, such as synaptic plasticity rules.
Related terms: Hebbian learning, spike-timing dependent plasticity, adaptive networks, lifelong learning.
Explanation: These systems adjust connection strengths based on activity patterns, enabling continuous learning and adaptation without catastrophic forgetting.
Example: A neural network that updates weights using a Hebbian rule when co-activation of units occurs.
Practical applications: Robotics that learn new tasks on the fly; AI personal assistants that evolve with user

preferences over years.
Challenges: Balancing stability with plasticity; scaling biologically plausible learning rules to large-scale architectures; validating performance against traditional backpropagation methods.

Multimodal Integration – the process of combining information from different sensory modalities into a unified percept.
Related terms: cross-modal binding, sensory convergence, audiovisual integration, Bayesian fusion.
Explanation: The brain integrates inputs such as sight and sound to improve accuracy and speed of perception, often weighting modalities according to reliability.
Example: The McGurk effect, where mismatched auditory and visual speech cues produce a fused perception.
Practical applications: Designing virtual assistants that combine voice and gesture cues for more natural interaction; developing autonomous vehicles that fuse LIDAR, radar, and camera