References

Harmon, L. W., Hansen, J. C., Borgen, F. H., & Hammer, A. L. (1994). *Strong interest inventory.* Stanford: Stanford University Press.

Holland, J. L. (1985). *Making vocational choices: A theory of vocational personalities and work environments* (2nd ed.). Englewood Cliffs: Prentice-Hall.

Parsons, F. (1909). *Choosing a vocation.* Boston: Houghton Mifflin.

Rounds, J., & Tracey, T. J. (1996). Cross-cultural structural equivalence of RIASEC models and measures. *Journal of Counseling Psychology, 43*, 310–329.

Savickas, M. L. (1999). The psychology of interests. In M. L. Savickas & A. R. Spokane (Eds.), *Vocational interest: Meaning, measurement, and counseling use* (pp. 19–56). Palo Alto: Davies-Black.

Tracey, T. J. G., & Ward, C. C. (1998). The structure of children's interests and competence perceptions. *Journal of Counseling Psychology, 45*, 290–303.

Stabilize

► Memory Consolidation and Reconsolidation

Stages of Internalization

Phases (or steps) of transformation of initially external action mediated by material or materialized tools into a mental (ideal) plan. Galperin introduced six stages of internalization: (1) formation of a motivation base of action; (2) formation of an orientation base of action; (3) formation of the material (materialized) form of action; (4) formation of the external socialized verbal form of action (overt speech); (5) formation of the internal verbal form of action (covert speech); (6) formation of the mental action; final changes, automatization, and synchronization of the action.

Star-Learning

► AQ Learning

State Anxiety

► Effects of Anxiety on Affective Learning

State-Dependency

► Mood-Dependent Learning

Statement

► Communication Theory

Static Mapping

► Initial State Learning

Statistical Learning in Perception

Nicholas B. Turk-Browne
Department of Psychology, Princeton University, Princeton, NJ, USA

Synonyms

Associative learning; Contextual cueing; Segmentation; Unsupervised learning

Definition

In cognitive psychology and cognitive neuroscience, statistical learning (SL) refers to the extraction of regularities in how features and objects co-occur in the environment over space and time. Such learning may be important for detecting and representing higher-order units of perception, such as words, scenes, and events. SL is defined by three criteria: First, it can operate over undifferentiated input, where only spatial and temporal probabilities can be used to determine which parts of the environment go together; other segmentation cues, such as grouping, are not required. Second, SL occurs incidentally as a by-product of perception, without intentional effort or conscious awareness. Third, SL is concerned with extracting how *particular* features and objects co-occur, resulting in knowledge about relationships between specific stimuli. These properties make SL well suited to the continuous and noisy sensory input we receive from the world.

Theoretical Background

Perception is concerned with interpreting information conveyed by our senses about the external environment, whether listening to music, enjoying a good meal, or recognizing a friend's face. Because we live in natural environments that have been stable for thousands of years, and we tend to inhabit the same artificial environments for years at a time, our perceptual systems are repeatedly confronted with very similar sensory information. Over time, repeated aspects of the environment – *regularities* – tune perception to the types of information most frequently encountered. Such learning may be critical for handling the deluge of sensory information that we experience from moment to moment, allowing for faster and more veridical perception of features and objects that appear briefly (e.g., due to eye movements or motion), or under variable or degraded conditions (e.g., due to changes in lighting or occlusion by other objects).

There are many types of regularities. Nature contains regularities in the layout of information (e.g., landscapes have horizontal horizons) and in the laws governing physical interactions (e.g., detached objects must be supported from below). Such natural regularities have altered perceptual systems over evolutionary time. For example, neurons in primary visual cortex are tuned to the properties of images regularly encountered by our species. Regularities also exist in terms of which *types* of information appear together. For example, academic robes are worn at commencements but not at the beach. Such semantic regularities are learned during our lifetime, and facilitate the recognition of objects that appear in appropriate contexts. Finally, regularities exist in terms of which specific *tokens* of information appear together. For example, a consistent sequence of landmarks is passed when navigating to a location: "the restaurant is after the cherry blossom tree, the independent bookstore, the busker playing saxophone, on the left". Importantly, such groupings are arbitrary (no semantic relationship between the objects), and the resulting knowledge is specific to the learned exemplars (a different tree, store, or person would not be helpful). The focus of this entry is on SL of regularities of this latter kind because: they are prevalent, they can be learned quickly in a laboratory setting, and learning is uncontaminated by preexisting knowledge of natural and semantic regularities.

Research on SL has roots in language acquisition. Early work focused on whether the boundaries between morphemes (smallest unit of meaning in a language) could be recovered from how phonemes (basic sounds in a language) are distributed/co-occur over time. The use of such regularities in language acquisition was first empirically tested in developmental psychology (Saffran et al. 1996). In particular, the speech input we get as listeners does not have reliable markers for where words start and stop, and so a key challenge for infants is not only to learn words and their meanings, but to locate the words in a speech stream in the first place. To test whether infants rely on statistical regularities to find word boundaries, a fake language can be constructed from arbitrary three-syllable words, such as *tupiro*, *golabu*, *bidaku*, *padoti*, which are then combined into a speech stream without pauses between words (e.g., *bidakupadotigolabubidakutupiro*...). Without such segmentation cues, learning of words requires extracting the transitional probabilities between syllables, that is, detecting that transitions within a word (*da* → *ku*) have higher probabilities than those spanning two words (*ku* → *pa*). After two minutes of exposure, infants can tell actual words apart from non-words – syllable sequences that did not occur during listening (e.g., *tulaku*) or occurred with lower probability (e.g., *dakupa*). Importantly, each individual syllable in the words and non-words is equally familiar, yet the infants expressed additional familiarity with the words. This study was used to argue that infants have powerful learning abilities, and that language acquisition may be partly experience-dependent.

Just as words can be defined by regularities of syllables, there is an analogous "language" of vision: Meaningful units of visual perception, such as scenes and events, can be defined by regularities of objects. Learning of visual regularities has been investigated for both temporal sequences and spatial configurations (e.g., Fiser and Aslin 2001). Temporal sequences arise in vision because events and actions (e.g., catching a flight) are defined by a specific progression of information (e.g., parking garage, ticket agent, metal detector, gate, seat), and because our eyes get information from one part of space at a time – resulting in a temporal sequence as we move our eyes around. SL of temporal sequences in vision is tested by showing a continuous stream of nonsense shapes (e.g., each letter is a shape: ABCGHIDEFABCJKL. . .) that, without observers' knowledge, has been constructed from shape triplets (e.g., ABC, DEF, GHI, JKL). After a few minutes of exposure to the stream, familiarity is higher for triplets (e.g., ABC) than for recombinations of the same shapes (e.g., AEI). Again, since each individual shape is equally familiar, additional familiarity for triplets indicates statistical learning of the shape transitions. Spatial configurations (e.g., the layout of objects in a room) contain regularities in the relative positions of objects (e.g., a lamp appears next to the sofa), and learning of such regularities may facilitate visual search and spatial navigation. SL of spatial configurations in vision is tested by showing grids of multiple shapes that, without observers' knowledge, have been constructed from shape pairs (if one member of a pair appears on the grid, its associate also appears in a fixed relation; e.g., to its left). Since many shapes appear in each grid, pairs can only be learned by extracting probabilities. After brief exposure to the grids, familiarity is higher for pairs than for recombinations of the same shapes into new arrangements, providing evidence of SL.

SL is ubiquitous in perception, occurring: (1) in auditory, visual, and tactile modalities; (2) for a range of statistics, including frequency and conditional probability; (3) throughout human development; (4) in other species, such as rats and monkeys; and (5) for many stimuli, from lower-level features to higher-level objects, scenes, and actions. The broad scope of SL suggests that it is a core cognitive function.

Important Scientific Research and Open Questions

Beyond demonstrating the scope of SL, research has begun uncovering the underlying nature of this process. For example, SL can occur without conscious intent or awareness, but is constrained by selective attention; SL involves memory systems linked to other types of implicit learning, such as artificial grammar learning and motor sequence learning; and SL happens quickly, after only a handful of repetitions of a regularity. This is a relatively nascent field of research, and thus three areas of active study are highlighted here.

First, while initial studies of SL used a single set of auditory or visual regularities, the environment contains multiple concurrent sets of regularities. For example, words in multiple languages, or the layouts of different cities. At a more basic level, sensory information arrives at the same time from multiple modalities (e.g., faces and voices), and about different feature types within the same modality (e.g., shape and color). It is important to assess how SL copes with this complexity to determine whether it is suited to more ecological settings outside of the laboratory. Studies have examined SL of regularities appearing simultaneously in different modalities or features, finding that if modalities/features are correlated (e.g., shapes appear in reliable colors) SL proceeds over multi-modal/multi-feature objects. In addition to appearing simultaneously, regularities can also appear sequentially, such as when learning a second language after already knowing a first. When presented with a speech stream constructed from one set of three-syllable words, followed immediately by a second stream constructed from a recombination of the same syllables into new words, learning occurs only for the *first* language (Gebhart et al. 2009) – in essence, SL of one set of regularities blocks learning of a second set. When two languages are separated in time by a pause, learning occurs for both languages. This raises an interesting possibility that only one set of regularities can be learned in a given "context," but that pauses and other contextual shifts may allow SL to restart. What counts as a context, and how much evidence about new regularities is required to overcome previously acquired knowledge remain open questions.

Second, most studies of SL test learning with the exact same regularities that were trained. For example, if presented with a triplet of shapes ABC during the initial learning phase, the test phase would examine familiarity for ABC. Without making changes between learning and test, it is unclear what has actually been learned. On one hand, SL could be very specific, resulting in knowledge that A is followed 1 s later by

B and then 1 s later by C. Alternatively, SL could be more abstract, resulting in knowledge that A, B, and C "belong together." These possibilities can be tested using *transfer* logic: if regularities can be recognized despite a change between learning and test, then the changed aspect is not an important part of what has been learned. Some changes do not matter: SL readily transfers to new temporal orders (training on ABC, testing on CBA), and even after the removal of all temporal information (training on ABC presented sequentially, testing on ABC presented simultaneously), suggesting that SL produces abstract knowledge that generalizes flexibly (Turk-Browne and Scholl 2009). Other changes do matter: after training on large spatial groups (shape triplets or quadruplets), SL does not transfer to the embedded subcomponents (constituent shape pairs), suggesting that SL of spatial configurations operates hierarchically at the largest spatial scale (Orbán et al. 2008). Characterizing what is learned during auditory SL remains an open and important question. For example, the transfer of visual SL to new temporal orders would be disastrous for word learning – reversing the syllables in any word changes its meaning, and most often produces a nonsense word.

Third, what purpose does SL serve beyond increasing familiarity? For auditory perception, segmentation of otherwise fluent speech during language acquisition could be a valuable consequence of SL. Segmentation is also important in visual perception, both for extracting edges, surfaces, and objects from the background, and for finding boundaries between events and actions – and again SL may be helpful. Another fundamental consequence of SL for perception is that the brain can use knowledge about regularities to *anticipate* what will appear next in the environment, allowing for more efficient processing when the predicted sensory information arrives. Objects that reliably predict what will appear next in a temporal sequence activate the anterior hippocampus and prime regions of visual cortex that are selective for the category of the predicted object. For example, faces that are predictive of scenes elicit anticipatory activity in scene-selective parahippocampal cortex (Turk-Browne et al. 2010). While the hippocampus is involved in explicit types of future-oriented behavior, such as planning and prospection, and category-selective visual regions are activated by effortful imagery, SL occurs without conscious awareness and thus triggers *implicit* perceptual anticipation. What consequences SL has for other cognitive processes, such as attention and working memory, remains an open and actively studied question. In sum, there are reciprocal interactions between SL and perception: SL occurs incidentally as a result of perception; in turn, learned regularities facilitate perception by allowing the brain to anticipate the future.

Cross-References

► Attention and Implicit Learning
► Exposure-Based Perceptual Learning
► Implicit Sequence Learning
► Incidental Learning
► Language Acquisition and Development
► Relational Learning
► Spatial Learning
► Task-Irrelevant Perceptual Learning
► Unconscious Learning
► Visual Perceptual Learning

References

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science, 12*, 499–504.

Gebhart, A. L., Aslin, R. N., & Newport, E. L. (2009). Changing structures in midstream: Learning along the statistical garden path. *Cognitive Science, 33*, 1087–1116.

Orbán, G., Fiser, J., Aslin, R. N., & Lengyel, M. (2008). Bayesian learning of visual chunks by human observers. *Proceedings of the National Academy of Sciences, 105*, 2745–2750.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*, 1926–1928.

Turk-Browne, N. B., & Scholl, B. J. (2009). Flexible visual statistical learning: Transfer across space and time. *Journal of Experimental Psychology. Human Perception and Performance, 35*, 195–202.

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *Journal of Neuroscience, 30*, 11177–11187.

Statistical Learning Theory

► PAC Learning