

Statistical Learning Subserves a Higher Purpose: Novelty Detection in an Information Foraging System

Ram Frost*, The Hebrew University; BCBL: Basque Center for Cognition
Brain and Language

Louisa Bogaerts, Ghent University

Arthur G. Samuel, Stony Brook University; BCBL: Basque Center for
Cognition Brain and Language; Ikerbasque: Basque Foundation for Science

James S. Magnuson, University of Connecticut; BCBL: Basque Center for
Cognition Brain and Language; Ikerbasque: Basque Foundation for Science

Lori L. Holt, The University of Texas at Austin

Morten H. Christiansen, Cornell University; Aarhus University

*Correspondence: ram.frost@mail.huji.ac.il

Authors' note:

This paper was supported by the Israel Science Foundation (ISF) Grant, Project 705/20, awarded to R.F. R.F. and M.H.C. were further supported by a Binational Science Foundation (BSF) grant number 2022082. L.B. received funding within the framework of the Odysseus programme from the Research Foundation Flanders (FWO), Project G0F3121N. The Spanish Ministry of Science and Innovation provided support to A.G.S. via Grant PID2020-113348GB-I00, and to J.S.M. via Grant PID2020-119131GB-I00. A.G.S. and J.S.M. were supported by the Basque Government through the BERC 2018-2021 and BERC 2022-2025 programs, and by the Spanish State Research Agency through BCBL Severo Ochoa excellence accreditation SEV-2015-0490 and CEX2020-001010-S. J.S.M. was also supported in part by U.S. National Science Foundation grant 2043903. LLH was supported in part by a US National Science Foundation grants 2346989 and 2420979 and a US National Institutes of Health grant R01DC017734. This work was not preregistered. Some of its theoretical claims were presented by Ram Frost at the 10th International Workshop on Advanced Learning Sciences (IWALS, Hanoi, December, 2024). We are indebted to Gershon Ben Shakhbar, Noam Siegelman, Shir Atzil, and Yuval Hart for their insightful comments on the manuscript.

Abstract

Statistical learning (SL) is typically assumed to be a core mechanism by which organisms learn covarying structures and recurrent patterns in the environment, with the main purpose of facilitating processing of expected events. Within this theoretical framework, the environment is viewed as relatively stable, and SL ‘captures’ the regularities therein through implicit unsupervised learning by mere exposure. Focusing primarily on language—the domain in which SL theory has been most influential—we review evidence that the environment is far from fixed: it is dynamic, in continual flux, and learners are far from passive absorbers of regularities; they interact with their environments, thereby selecting and even altering the patterns they learn from. We therefore argue for an alternative cognitive architecture, where SL serves as a subcomponent of an *information foraging* (IF) system. IF aims to detect and assimilate *novel* recurrent patterns in the input that deviate from randomness, for which SL supplies a baseline. The broad implications of this viewpoint and their relevance to recent debates in cognitive neuroscience are discussed.

Keywords: Statistical learning, language, recurrent regularities, information foraging, novelty detection, curiosity.


Reappraising the role of statistical learning in cognition

Statistical learning (SL) has been taken to be a central mechanism by which organisms learn recurrent patterns in their environment (see Santolin & Saffran, 2018, for a review), from temporal regularities in audition (e.g., Saffran, Aslin, & Newport, 1996), to spatio-temporal regularities in vision (e.g., Fiser & Aslin, 2001; Shapiro & Turk-Browne, 2015) and haptics (Conway & Christiansen, 2005), cross-situational sound-meaning relations (e.g., Yu & Smith, 2007), implicit regularities in probabilistic classification (e.g., Knowlton, Squire, & Gluck, 1994), and cross-modal motor sequences (e.g., Kemény & Németh, 2017). Research on SL typically assumes that learning of such regularities in sensory input is implicit and unsupervised, occurring with mere exposure, not requiring feedback (e.g., Santolin & Saffran, 2018), in contrast to supervised learning in which reward feedback drives changes in behavior¹. Across this extensive literature, the role of SL is to detect and assimilate the covarying structures in the environment, with the main purpose of facilitating the processing of likely events in the upcoming sensory input (e.g., Turk-Browne, Jungé, & Scholl, 2005; Frost, Armstrong, Siegelman, & Christiansen, 2015; Yan, Ehinger, Pérez-Bellido, Peelen, & de Lange, 2023). The assimilation of regular structures in time and space, therefore, is typically assumed to be the *goal* of SL. In the last three decades, an extensive body of research has focused on mapping how learners acquire such statistical knowledge (e.g., Mareschal & French, 2017; Schapiro, Turk-Browne, Botvinick, & Norman, 2017; Siegelman, Bogaerts, Armstrong, & Frost, 2019; Henin et al., 2021; see Frost, Armstrong, & Christiansen, 2019 for discussion).

Here we argue for an alternative cognitive architecture, in which SL instead is a subcomponent (albeit a key subcomponent) of an *information foraging* (IF) system. The purpose of the IF system is to detect *changes* in the environment, for which SL provides continually updated baselines. The proposed cognitive architecture is schematically portrayed in Figure 1.

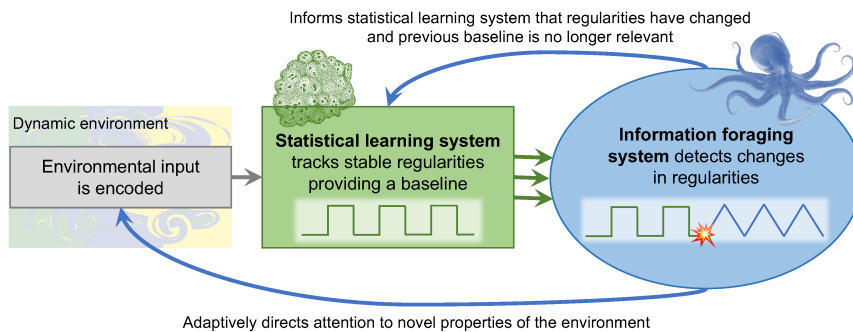
¹ Some neural network models—e.g., SRNs (Elman, 1990) and Large Language Models [LLMs] such as GPT-3.0 (Brown et al., 2020)—learn statistical patterns in language through *self*-supervised learning, where the networks make predictions for what should come next and use the subsequent actual input as indirect feedback on the correctness of those predictions, adjusting their weights accordingly (see Contreras Kallens et al., 2023, for discussion). Thus, self-supervised learning is arguably a form of unsupervised learning, and could be a plausible candidate mechanism supporting SL. Interestingly, LLMs are able to reproduce or emulate human language more closely when they also receive supervised learning from human feedback in addition to self-supervised learning (as in GPT-Instruct and ChatGPT; Ouyang et al., 2022).

A. The Statistical Learning View



A sea sponge as metaphor for statistical learning. The learner is bathed in information and absorbs regularities incidentally and passively.

B. The Information Foraging View



An octopus as metaphor for information foraging. The learner of regularities explores, interacts with, and alters the environment.

Figure 1: Cognitive architectures according to the SL view (panel A) versus IF view (panel B). On the conventional SL view, the focus is on implicit learning of regularities of a relatively stable environment. On the IF view, the SL learning system subserves the IF system by providing a baseline, and is informed by it regarding changes in regularities. Following informative changes in regularities, the IF system diverts attention to novel dimensions of the environment; that is, the focus of IF is change or deviations from prior regularities, rather than regularity itself. The second crucial aspect of IF is that while much learning may be implicit, the IF learner actively interacts with its environment, selecting or even altering available information.

We should emphasize at the outset of our discussion that there is no formal agreed-upon definition of what SL phenomena include and, therefore, what computational mechanisms they involve. Some approaches argue that all (or almost all) learning is, in fact, statistical (see Frost et al., 2019, for review). Hence, to be clear, throughout this paper we use the term “SL” to specifically denote *the well-researched robust learning mechanisms that allow organisms to implicitly and unintentionally assimilate recurrent patterns in their sensory input by mere exposure, in an unsupervised manner, a mechanism which has been taken to be a central device in mirroring the regularities of the environment.* An implicit tenet underlying this approach is that the environment is, to a large extent, set and stable in terms of its structure, allowing organisms to map its structural regularities through an implicit SL process.

The IF approach builds on a fundamentally different view of the environment. Its starting point is that across domains, statistical regularities in sensory input are typically dynamic and continuously

changing. Hence, the primary goal of the system is to detect these changes. The IF system thus aims to detect *novel* recurrent patterns in the input that deviate from randomness and from past baseline regularities, in the service of efficient and continuously adaptive processing. In support of this proposal, we draw primarily on examples from language because this is the domain for which SL theory has been most influential, but as we discuss further on, the motivation and principles we describe apply more generally.

The IF approach to cognitive architecture we put forward departs from the conventional SL view across a range of important dimensions, as summarized in Table 1, which also provides a roadmap for this article. In the following, we explicate each of these dimensions in detail, provide behavioral and neurobiological evidence in support of the proposed architecture, discuss how the IF approach aligns with important higher cognitive functions such as curiosity and creativity, and outline a blueprint for a novel research agenda to investigate information foraging.

	The Statistical Learning View	The Information Foraging View
Role in cognitive system	SL stands alone	SL is a subcomponent of IF, providing a baseline for detecting changes.
Nature of the environment	Relatively stable regularities	Dynamic, structural properties in flux
Nature of the learner	Passive, absorbing regularities via implicit passive learning	Active, exploring, interacting with and altering the environment
Aim of the system	Assimilate regularities	Information foraging
Target of perception	Expected events	Unexpected events
End state	When regularities are assimilated	None; continuous foraging
Selectivity	Nonselective learning regardless of information value	Selective learning of informative regularities
Neurobiological characteristics	Increased activation with increased regularity	Indifferent to high regularity and to randomness
Individual differences	Sensitivity to regularity	Sensitivity to non-random changes in regularity

Table 1: *Main contrasting features of SL and IF approaches to cognitive architecture*

The ever-changing nature of the linguistic environment

While the building blocks of a given language and their basic patterns of co-occurrence in speech, sign, or print are typically characterized as well-defined and rigid, their actual realization in conversational context is messy, variable, and muddled, requiring constant adaptation and

accommodation. While there are of course stable aspects to the linguistic environment at abstract levels of description, deviations from central tendencies constantly occur at the level of the actual input. These deviations signal high informational value. Efficient language processing thus requires that the constant changes in regularities be perceived rapidly, allowing effective comprehension of the novel linguistic input. This can be easily demonstrated in the domains of speech, print, or sentence processing.

In spoken language processing, the inventory of speech sounds comprising a language can be relatively well-defined, but the statistical regularities in speech streams are characterized by varying phonetic realizations. Speakers' productions diverge due to variation in physical characteristics (sex, age, size), social characteristics (gender, social identity), dialect, accent, speech rate, vocabulary, and cultural backgrounds. Listeners must not only constantly accommodate this variation to achieve phonetic constancy (see Luthra, 2023, for a recent overview), but also simultaneously leverage the information it provides for purposes ranging from talker identification (e.g., Perrachione, Pierrehumbert, & Wong, 2009), to learning talker-specific phonetic idiosyncrasies that improve later perception (e.g., Norris, Cutler & McQueen, 2003; Nygaard & Pisoni, 1998), to making a surprising array of physical and social inferences (e.g., Krauss, Freyberg, & Morsella, 2002; Munson & Babel, 2007).

The predictability of word sequences significantly varies across different age groups and in different linguistic environments. In everyday conversations, speakers shift swiftly between different registers and codes, adapting their style of talking to their interlocutors, whether that be authority figures, like a police officer or our boss, or more affiliative dialogic partners, like parents, children, or friends (see Goulart et al., 2020). In such different linguistic environments, interlocutors' language further tends to rapidly adapt to context-specific parameters of the interaction at every level, potentially mitigating those sources of variability (Garrod & Pickering, 2004). In many cases, conversational partners' speech becomes more similar in terms of articulatory, acoustic, and prosodic details (e.g., Kim, Horton, & Bradlow, 2011; Lee et al., 2018; Pardo, 2006), syntax (e.g., Bock, 1986), semantics (e.g., Dideriksen, Christiansen, Tylén, Dingemanse, & Fusaroli, 2023), and kinetic alignment of head and hands (Trujillo, Dideriksen, Tylén, Christiansen, & Fusaroli, 2023). But in other cases, interlocutors will deviate from one another to provide new information when it is helpful to solve a particular task

(e.g., Didriksen et al., 2023; Trujillo et al., 2023; see Fusaroli, Rączaszek-Leonardi, & Tylén, 2014, for discussion). Speakers and listeners appear to attune exquisitely to novel information in conversation, with speakers producing words referring to new information for the first time with greater clarity than words referring to given (old) information (e.g., Fowler & Housum, 1987). Speakers (and writers) also appear to strive to maintain ‘uniform information density’ over time, via word choices and prosodic structure (e.g., Aylett & Turk, 2004; Frank & Jaeger, 2008; Genzel & Charniak, 2002; Gibson et al., 2019), even to the level of discourse properties (Asr & Demberg, 2015). Speakers and listeners are sensitive to dynamic changes in discourse-relevant semantic and phonological neighborhoods even as they alter those neighborhoods themselves (e.g., by settling on a shared vocabulary in a novel task, as in Brown-Schmidt & Tanenhaus, 2008—see Brown-Schmidt et al., 2015, for a review). Hence, the statistical properties of the speech environment are anything but fixed or stable, and the main game in conversational context is to perceive and adapt to the ever-changing novel structure — not to simply assimilate stable aspects of the environment.

A similar state of affairs characterizes reading. Predictability of letters within words is to a large extent stable, constrained by orthographic and phonotactic rules (e.g., Siegelman, Kearns, & Rueckl, 2020), and linguistic form enables some general predictions at higher levels of abstraction (e.g., Snell & Theeuwes, 2020). However, beyond this, readers are constantly faced with novel regularities. The predictability of printed words, which drives ocular movements during text reading, changes significantly as a function of variation in writing style, the period when the text was written, and the type of text being processed, from movie subtitles to newspapers. Efficient processing thus requires the reading system to detect and adjust to such changes as rapidly as possible. Recent research tracking eye-movements indeed shows that readers rapidly perceive and adapt to specific syntactic structures characteristic of the writing style (Yan & Jaeger, 2020), and to expected sequences of word-lengths from sentence onset in a given text, to optimize ocular movements (Snell & Theeuwes, 2020). For example, when presented with sentences of uniform word-length, readers adjust their preferred saccade length incredibly rapidly; just a few exemplars of a given word-length suffice (Cutter, Dreighe, & Liversedge 2017, 2018). Such context-dependent adjustments are well documented for speech as well. If listeners hear speech segments that are intentionally made ambiguous, with the lexical context

providing disambiguation, they rapidly recalibrate their specification of the segments accordingly (e.g., Norris, McQueen, & Cutler, 2003). Language users also attune rapidly to changes in phonotactics when they produce speech (Dell, Reed, Adams, & Meyer, 2000), or simply listen to it (e.g., in the context of a lexical decision task; Onishi, Chambers, & Fisher, 2002).

In the same vein, syntactic structures in spoken and written language are anything but uniform, reflecting the immense creativity characterizing human linguistic interaction (Christiansen & Chater, 2022). These distributional changes have a direct impact on online sentence processing (Wells, Christiansen, Race, Acheson, & MacDonald, 2009). In addition, readers must also contend with the different distributions of syntactic regularities associated with different genres of writing (Snell & Theeuwes, 2020), from academic treatises and newspaper articles to fictional books and blog postings (Goulart et al., 2020).

To summarize, while linguistic input allows for predictions at various levels of abstraction, speakers, listeners, and readers must constantly adapt to novel, ever-changing structures in the input stream, rather than merely encoding stable ones.

The active nature of the learner

The main appeal of SL is in the robustness and power of implicit learning mechanisms, which are already operating in newborns (e.g., Bulf, Johnson, & Valenza, 2011), and which do not require overt attention (Saffran et al., 1997). In the latter seminal study of Saffran and colleagues, children showed learning of transitional probabilities of speech sounds heard in the background while they drew pictures. In typical experimental studies of SL, learners are not informed about the existence of statistical regularities, nor warned of a subsequent test of their knowledge of them. The exposure phase involves passive listening or viewing of continuous input streams, yet participants (on the average) consistently show learning of the recurrent patterns in the input. These results paint the statistical learner as an efficient passive absorber of environmental regularities, which are assimilated via robust implicit learning mechanisms. In Figure 1, this view is reflected in the “sea sponge metaphor”: the learner is immersed in statistical regularities and assimilates them (see Tandoc et al., 2024, for a similar description).

The IF perspective does not dispute the existence of implicit unsupervised learning of recurrent patterns; indeed, such learning has been demonstrated from humans and other primates to songbirds (e.g., Santolin & Saffran, 2018; Lu & Vicario, 2014), and serves as a key subcomponent of IF. Crucially, though, the focus of IF is on the learner as an active explorer, an information forager, who registers recurrent patterns, but actively intervenes in regularity learning by exploring, interacting with, and altering its environment. In Figure 1, this view is depicted in the octopus metaphor, in contrast to the passive sea sponge. Regularities inform the learner where to forage — *away* from highly predictable patterns and away from randomness related to simple noise (i.e., the environment is not only dynamic but also noisy), allowing it to direct attention and action to times, places, and non-spurious events with potentially high informational content or reward. Deviations from regularity are one important source (among many) that guide perception, attention, and action in service of adaptive learning. Note that while the sponge is a metaphor for the system(s) comprising passive SL through mere exposure, the octopus is a metaphor for both SL system(s) and the mechanisms supporting IF. These comprise the organism itself — its body and perceptual/cognitive systems can be actively and intentionally directed towards information seeking and learning.

Our take is that regularity learning, like all aspects of cognition, is inextricably linked to perception, action, and the environment (Sheya & Smith, 2019), as well as neural and genetic activity (Gottlieb, 2007; Gottlieb & Oudeyer, 2018). Theoretical perspectives such as dynamical system approaches to development (e.g., Smith & Thelen, 2003) and probabilistic epigenesis (Gottlieb, 2007), emphasize bidirectional interactions among all these levels. Hence, a learner is not simply shaped by the environment and its regularities. An active learner alters the environment through their actions, whether in speech communication as we have described above, or, say, learning words or linking them to objects. For example, as Smith, Yu and colleagues have documented, an infant’s actions modulate the multimodal context for language learning. When infants hold and manipulate an object, they change the visual and haptic context, and such actions, as well as where they direct attention in the visual world, can also modulate the language produced by adults around them (e.g., Slone, Abney, Smith, & Yu, 2023; Suanda, Barnhart, Smith, & Yu, 2018; Suanda, Smith, & Yu, 2017). Similar interactions with people around them influence other domains of learning (e.g., Karmazyn-Raz & Smith, 2023; Smith,

Jayaraman, Clerkin, & Yu, 2018; Suarez-Rivera, Smith, & Yu, 2019; Yu, Zhang, Slone, & Smith, 2021).

How the learning of regularities is contingent on active responses has been recently shown in the *singleton paradigm*, where participants search for a shape singleton target and are asked to respond to it. Li, van Moorselaar, and Theeuwes (2024) reported that if a target's location is predicted by the location of a target in the preceding trial, the execution of an arbitrary key response for both trials of the pair is needed for learning the across-trial statistical regularity. Passively attending to target locations did not result in any learning of the target's spatial contingencies.

To clarify, the active nature of the learner does not imply conscious awareness or meta-awareness. It implies that learners do not automatically absorb the recurrent patterns in sensory input, but continuously interact with their environment and change it, whether consciously or not. The earliest ways in which infants shape their environments are subtle and likely occur in the absence of any awareness (when an object or entity in the visual world captures an infant's attention, adults are more likely to speak about it). However, over experience, infants increasingly leverage this to intentionally shape adults' behavior by directing their gaze or body (by pointing, approaching, or handling) selectively at aspects of the world they find interesting, especially deviations from regularity – including deviations they themselves cause.

Aim of the system

Given the ever-changing environment, the IF approach to cognitive architecture assumes that the priority of cognitive systems lies in continuously detecting *novel* co-occurrences and other coherent covariations in the input. It regards organisms as “information foragers” (see Pirolli & Card, 1999, for coining this concept), where “information” means that a meaningful change in patterns of co-occurrences in sensory input has occurred. Our approach in this context is related in part to the information-theoretic notion of *information* (Shannon, 1948), which ties the informational load that events carry to the inverse of their predictability. By this view, highly predictable events carry very little information. However, in the context of regularity learning, random events are also uninformative. Hence, an information forager discounts random events as well as highly predictable events since both

carry little information. If the primary aim is to detect the novel regularities in the environment, this requires a mechanism that generates a reference against which changes in input regularities can be perceived. Conventionally, SL has been seen as a mechanism aiming to perceive and assimilate *stable* structural regularities in the environment. In the IF approach, SL mechanisms that track statistical regularities have a different core purpose: they provide the baseline from which a change in co-occurrence can be detected. To be clear, we consider the detection of novel regularities and the generation of baseline references as *functionally* distinct at the cognitive level of description. As discussed further below, extant data raise the question of whether they can be implemented within one computationally and neurobiologically unified mechanism.

Consider a typical visual SL task that presents a continuous stream of shapes or artificial letters appearing in triplets, where elements within triplets are fully predictable, and elements following triplet boundaries are less predictable (e.g., Fiser & Aslin, 2002; Turk-Browne et al., 2005; Saffran et al., 1996; Siegelman & Frost, 2015). Conventional SL theory assumes that learning serves to facilitate processing of the predictable stimuli (see for example, Turk-Browne, Scholl, Johnson, & Chun, 2010; Siegelman, Bogaerts, Kronenfeld, & Frost, 2018). Some recent findings, however, suggest that learners in such tasks track novelty. For example, a study that tracked EEG activity during the continuous presentation of triplets of visual stimuli revealed that increased pattern repetitions resulted in increased beta-band activity, which has been associated with sensory prediction (e.g., Arnal & Giraud, 2012) and top-down modulation (e.g., Hipp, Engel, & Siegel, 2011). Importantly, however, this top-down modulation was present *at triplet transitions*, where a novel shape is about to appear, and not within triplets, where shapes are predictable (Bogaerts, Richter, Landau, & Frost, 2020). As we further discuss in detail, the suggestion that probabilistic knowledge can upweight surprising rather than predictable events, favoring novelty over familiarity, is now acknowledged across different domains of cognitive neuroscience.

The target of perception

SL theory assumes that learning regularities in the environment enables their exploitation by facilitating perception of and action upon expected events. As such it considers likely recurrent events

to be the foremost target of perception and learning. In contrast, an IF approach holds that while predictable events can be favored by various top-down mechanisms, *unexpected* patterns of events constitute informational novelty, and are, therefore, the main target of perception. Indeed, as early as Pavlov (1927), orienting behavior has been taken as a primary mechanism aimed at detecting the slightest change in the environment. Further, it was argued that recurrent presentations of stimuli result in neuronal representations that encapsulate the stimuli's specific features, so that all sensory input could be compared with the existing neuronal models, and a mismatch between novel input and the models would result in an orienting reaction (Sokolov, 1963). The competition between baseline habituation and novelty has been shown to drive orienting behavior and foraging for visual information (e.g., Sirois & Mareschal, 2004). While there is extensive variability in the definition of what constitutes “novelty” in this context (see, e.g., Gati & Ben-Shakhar, 1990, for discussion), the view that neuronal models of recurrent and expected events serve as a baseline to flag novel and unexpected recurrent events converges with our view (and see, e.g., Egner et al., 2010; Meyer & Olson, 2010; Kumar, Kaposvari, & Vogels, 2017; Richter et al., 2018, for evidence of surprisal detection in the neural domain).

Recent work on attention provides additional support for the IF approach. The *learned distractor suppression* literature shows that when the distractors in a series of search displays frequently occur in the same location, they capture attention significantly less (see Theeuwes, Bogaerts, & van Moorselaar, 2022 for review). These results suggest that attention is modulated by distributional regularities in the environment, prioritizing novelty. Only distractors occurring at *unexpected* locations compete strongly for attention. Learned suppression of predictable distractors has also been observed for distractor features. For example, a distractor in a specific color loses its ability to capture attention with repeated exposure (Vatterott & Vecera, 2012; see Geng, Won, & Carlisle, 2019, for review). In addition, if distractors are highly frequent this can eliminate capture, even if their location and features are unpredictable (Bogaerts, van Moorselaar, & Theeuwes, 2022; Won, Kosoyan, & Geng, 2019). Consistent with our approach, it seems that when distractors become part of the baseline, and are no longer novel or surprising, they become less salient. Similar effects may be at play in the speech domain. Understanding what a stranger has said is easier in a ‘cocktail party’ scenario when simultaneous,

competing speech is produced by a highly familiar talker (one's spouse) vs. when the competing voice is that of a novel talker (Johnsrude et al., 2013). Moreover, frequently occurring sounds lead to reduced perception of similar sounds, an effect that has been extensively studied in the selective adaptation literature (Eimas & Corbit, 1973; see Samuel, 1986, for a review).

The system's end state

The successful assimilation of the structural regularities present in the environment is typically viewed as the SL system's end state. While SL naturally assumes that the learned regularities can be continuously updated given gradual changes, the view is that the system has done its job when it has learned the current regularities in a given domain. This view is well reflected in the experimental tasks devised to probe SL, which by and large, do not engage in change, perception of change, adaptability to change, and continuity (see Frost et al., 2019, for a review). Rather, participants are typically presented with an input sequence and then tested on whether they have learned the regularities that were embedded therein. Learning is inferred when there is evidence that participants have assimilated the structural regularity embedded in the input stream, whether visual or auditory (e.g., Siegelman, Bogaerts, Christiansen, & Frost, 2017). While, admittedly, these tasks were initially designed to provide a proof of concept that the statistical regularities in the input stream can be learned, in our view, they implicitly entrenched theoretical approaches to SL to consider it as a system whose end state is to map the existing structure of a stable environment.

In contrast, IF assumes ongoing foraging for novel information, and thus that there is *no* end state. In the domain of language, given the dynamic changes in the input, IF requires the perceiver to be constantly adapting to what is different. For example, phoneme perception is immediately impacted upon encountering a new voice that shifts in mean acoustic spectra (Holt, 2005; Huang & Holt, 2012), or when hearing a foreign accent in which acoustic input dimensions differ in their correlation (Idemaru & Holt, 2011; Hodson, Shinn-Cunningham, & Holt, 2023). Indeed, as we discuss further below, given that most SL experiments involve stimuli that are often very novel, these experiments may be better construed as implicating, at least initially, sensitivity to new structure.

Of course, at the end of the day, organisms do represent knowledge about stable structural properties of the environment, language included, and this knowledge facilitates perception and action. From the IF perspective we are proposing, SL mechanisms are continuously at work, and if baseline regularities are recurrently registered across time, without much change, this information will be represented (subserving interaction with the environment; see, e.g., Schapiro et al., 2017) and updated to reflect gradual changes that may emerge over time. One could argue that, given this, IF subserves the SL system rather than the other way around. However, since the statistical regularities in sensory input are dynamic and continuously changing, adaptation to these novel changes is the system's primary goal. Importantly, how much knowledge of stable real-world linguistic regularities is acquired from mere exposure to input regularities through implicit and unsupervised SL mechanisms, and to what extent additional kinds of learning (e.g., supervised) might be necessary, remains to be determined (see Bröker et al., 2024, for review).

For example, concurring with our IF octopus metaphor, the alignment of statistical input regularities with active behavior amplifies learning above and beyond what is possible across passive exposure alone. An illustrative example comes from novel nonspeech sound clusters designed to mimic the statistical structure and complexity of English consonant categories. These categories are not acquired with passive exposure (e.g., Wade & Holt, 2005; Emberson, Liu & Zevin, 2013; Roark, Lehet, Dick, & Holt, 2022). Nonetheless, they are rapidly learned across the same time course when they are embedded in an unrelated active task structured such that the sounds are not essential to the task, but learning their structure supports success. This learning robustly generalizes to novel exemplars (Roark et al., 2022), alters cortical representations (Leech et al., 2019), and persists over days (Gabay, Karni, & Holt, 2023), even without knowledge that categories exist. The structure is discovered via its utility in supporting behavior. Active engagement in a rich, multimodal perceptual environment (typical of most natural behaviors) may encourage foraging for information that directs learners to specific statistical regularities among the essentially infinite informational contingencies that exist even in simple real-world environments (Roark et al. 2022). Thus, regularity that is difficult to extract across passive listening is readily learned, perhaps in a form of 'self-supervised' learning (Lim, Fiez, & Holt,

2019), by virtue of coarse alignment of statistically structured input with behaviorally relevant actions, objects, and events.

In this context, both endogenous generation of feedback as in self-supervised learning and exogenous (supervised) feedback fit with the view that the learner actively engages with the environment. We assume that the division of labor between the different learning mechanisms is likely to differ across domains (e.g., reading, syntax, second-language acquisition), and potentially across individuals. For example, explicit instruction and feedback play a significant role in reading acquisition (see Rastle, Lally, Davis, & Taylor, 2021), which is not the case for native spoken language acquisition.

Selectivity

In conventional SL theory and in typical SL research, the structural regularities in the environment are defined independently of whether they are informative to the learner and of the specific context of learning (e.g., Lelonekiewicz, Ktori, & Crepaldi, 2020; Fiser & Aslin, 2002). Indeed, implicit passive learning of regularities present in sensory input (whatever computational mechanisms are assumed) is taken to be nonselective, in the sense that the system is set to assimilate and absorb the statistical co-occurrences present in the input through mere exposure, whatever the regularity is, regardless of its informational value to the learner. While having a mechanism that registers recurrent regularities non-selectively might have some advantages in terms of simplicity, we argue that passive, non-selective SL mechanisms are limited in their usefulness for driving higher-order behavior. Instead, their importance comes from their role as a baseline-providing subcomponent of a higher-order IF system, which is selective. The IF system must be selective because any given environment presents the learner with a myriad of possible regularities, and it is the current context that determines which are “informative” and which are not for a given organism.

The importance of selectivity and the role of contexts has not gone completely unnoticed within the domain of SL research. For example, a good demonstration that learners are not passive absorbers of regularities in sensory input and that the informativeness of the signal modulates learning is children’s ability to assimilate regularities such as ABB (e.g., generalizing *le-di-di* to *ko-ga-ga*; see Marcus et al., 1999). Marcus, Fernandes, and Johnson (2007) have shown that when children hear ABB

in speech sounds, they learn the patterns, but when they hear ABB in non-speech sounds, such as sine-wave tones, they do not. Importantly, in a subsequent study, Ferguson and Lew-Williams (2016) demonstrated that if children were previously exposed to a video of two persons communicating in tones (a communicative context), learning does occur for tones, just as it does for speech sounds (see Saffran et al., 2007, for a related finding with images children found interesting [dogs] vs. not [shapes, used by Marcus et al., 2007]). This shows that there are preferences regarding what regularities should be attended to and what regularities can be ignored. The context determines what in the environment carries important information for a given species and what does not. In this example, regularities such as speech sounds that subservise communication within humans appear to be a primary filter for selection.

This tension between informative and non-informative regularities has also been acknowledged in the domain of perception with a similar argument; it is computationally infeasible to sample all available information in the very complex real-world environment (Braunlich & Love, 2022). Learning is selective because IF requires learners to eventually ignore regularities in the input that are uninformative (with respect to goals relevant for the organism to thrive in its niche), and focus on the informative ones. Note that making sense of which variations in the input are informative and which are not in itself requires learning. In fact, an essential part of development can be regarded as mastering this distinction. For example, very young infants are sensitive to subcategorical (subphonemic) variations in speech (allowing them to distinguish non-native speech sound contrasts, such as the /r/-/l/ distinction for Japanese infants), but during the first year of life they learn to divide acoustic-phonetic space in ways that are optimized for the language they are immersed in, and they lose the ability to distinguish most contrasts that are not relevant for that language (e.g., Werker, Gilbert, Humphrey, & Tees, 1981). In the domain of reading, proficiency has been shown to be related to the extent to which a reader relies on systematic (e.g., orthography to phonology) regularities versus spurious (e.g., arbitrary semantic cues such as imageability) regularities that are characteristic of the orthography (Siegelman et al., 2020).

Successful SL will result in lasting representations aligned with language regularities. This internal knowledge provides another form of selectivity for IF. In lexically-mediated perceptual learning (Norris et al., 2003), informativeness is seen in at least two ways. First, shifts in phoneme boundaries seem typically to be learned in a talker-specific way (segregating novel statistics experiences in speech from

a particular talker from baseline, rather than assuming the overall baseline for the language has changed; e.g., Eisner & McQueen, 2005). Second, learning is blocked when an alternative explanation is available for phoneme boundary shifts (e.g., seeing that the speaker has a pen in her mouth; Kraljic, Samuel, & Brennan, 2008). In dimension-based statistical learning in speech, individual differences in acoustic cue weighting (a reflection of long-term speech representations) predict how local speech input regularities shift speech categorization (Wu & Holt, 2022). Thus, statistical learning under control of information foraging is selective, can be adaptively conservative in its scope, and can involve foraging of internal representational space as well as local input.

Neurobiological characteristics

In the conventional view, SL serves to provide the organism with an internal mirror of external regularities in the environment, so that greater external regularity should lead to faster learning. From a neurobiological perspective, the assumption of this view is that the more regularity there is in the input, the more neural activity there should be in the medial temporal lobe (MTL) as well as in early sensory cortices, leading to faster assimilation of the external statistics (e.g., Schapiro et al., 2017). IF systems operate differently, because both full randomness and full regularity are uninformative, as they do not represent meaningful novelty. This view of a “Goldilocks” range of information concurs with behavioral findings showing that children direct attention to events that are neither too simple nor too complex (e.g., Kidd et al., 2012; see also Forest, Siegelman, & Finn, 2022, for adults). Hence, the neural architecture of an IF system should not be responsive to randomness, nor to completely predictable patterns. Indeed, neuroimaging studies have identified neural populations that track uncertainty non-monotonically following an inverted U-shaped function (with stronger response for moderately unexpected inputs, but low response for both highly expected and highly random inputs), in both the visual and auditory cortices (Nastase, Iacovella, & Hasson, 2014; Hasson, 2017). These neural systems do not respond to full randomness or full regularity as these are alike in terms of informativeness (or lack thereof); instead, they are tuned to the *moderate regularities* in the sensory input.

Extensive neuroscience work has investigated the interconnection of the reward system and salience network (Seeley et al., 2007) with the violation of expectation. For example, substantial

research has tied the amygdala, insula, and dorsal anterior cingulate cortex to computations flagging novelty and surprise (e.g., Kolling, Wittmann, Behrens, 2016; see Vassena, Holroyd, & Alexander, 2017, for review), thus potentially mediating information foraging. A different body of evidence comes from work on reinforcement learning and the dopaminergic system showing that novel and/or infrequent information is encoded by dopaminergic neurons in the striatum (e.g., Schultz, Dayan, & Montague, 1997). This research, however, has mainly focused on prediction errors regarding upcoming reward given changes in its probability (e.g., Behrens, Woolrich, Walton, & Rushworth 2007), and not on general deviation from baseline regularity in implicit unsupervised learning. Hence, the question of whether and to what extent IF is rewarding beyond paradigms of supervised learning requires further investigation.

A hint comes from neuroimaging of the learning of nonspeech categories, described above. Recall that these regularities are not learned with passive exposure but are robustly acquired when the regularity aligns with behavior in an active task (Wade & Holt, 2005; Roark et al., 2022; Gabay et al., 2023). Examining such learning with fMRI reveals that that posterior striatum (especially caudate and putamen) is sensitive to statistical regularity. When actions and events in an active task incidentally align with sound categories that are defined by well-structured statistical regularities, the posterior striatum is recruited to a greater degree than it is among participants who engage in the same task with statistically less-well-structured categories. The magnitude of striatal activation is associated with better behavioral learning outcomes. Thus, when statistical regularities align with actions and events in the environment, ‘self-supervised’ learning signals available via the posterior striatum may boost learning beyond what is possible through passive exposure alone (Lim, Fiez & Holt, 2019).

Along with our present IF approach, a recent framework tying learning and memory to curiosity (the Prediction, Appraisal, Curiosity, and Exploration [PACE] Framework; Gruber & Ranganath, 2019) demonstrates how, in general, prediction errors and detection of novelty increase attention and exploration through modulation of activity in dopaminergic circuits. Hence, while SL research has typically focused on the role of the MTL in registering regularities mainly through hippocampal activation (e.g., Schapiro et al., 2017), an IF approach would incorporate a larger set of neurobiological

mechanisms that simultaneously consider a) sensitivity to deviation from regularity, b) reward systems, and c) memory.

Individual differences

Both SL and IF accounts assume that there are individual differences in pattern sensitivity (see Frost et al., 2015; Siegelman et al., 2017, for discussion). Since the environment is not only dynamic but also noisy, the challenge to any learning system, SL and IF alike, is to distinguish meaningful patterns related to regularities, from those related to noise. Separating noise from signal requires something akin to a time window across which random noise will average out based on not having any predictive value, whereas regularity of the signal will remain. If the sampling window is too short, the system will change/reorient/relearn with every bit of fluctuating noise. If the time window is too wide, shifting patterns could go unnoticed. One possibility is that such a sampling window may be implemented in a literal fashion, as in a time series analysis. However, this may not be a likely solution because learners cannot hold on to multiple input patterns before processing them, and our limited memory abilities leaves little room for backtracking (cf. the Now-or-Never bottleneck, Christiansen & Chater, 2016). Instead, we lean toward a more metaphorical interpretation of the sampling window, such as what might be observed in recurrent networks trained on sequences. Here, there is no explicit sampling window, but coherent signals will triumph over noise because consistent patterns (relative to dynamic contexts) will be the primary driver of weight changes over time. Interestingly, work by Karuza et al. (2016) indeed suggests that the assimilation of structure leads to decreased environment sampling, causing learners to overlook pattern shifts and to display a bias toward their initial experiences (see also Bruner & Postman, 1949). We assume that individuals differ in this respect, that is, how they optimize their ‘sampling window’ in a given context.

Although research from an SL perspective targets individual sensitivity to detecting regularities (e.g., Misyak, Christiansen, & Tomblin, 2010, Siegelman & Frost, 2015), IF targets individual sensitivity to detecting *changes* in regularity. Given that the environment is typically dynamic rather than stable, individuals are expected to differ in their perceptual sensitivity to the ongoing dynamic changes in regularities in sensory input, and their efficiency in acting on these changes. This again

would be tied to differences in optimizing the sampling window in a given context. Whereas it is possible that sensitivity to stable regularities goes hand in hand with sensitivity to change, it is also possible that these would be two dissociable dimensions of inter-individual variance. Note that IF also assumes substantial individual differences in sensitivity to informativeness. These differences are found not only across development (efficient IF requires the ongoing learning of which regularities to assimilate and which to ignore), but also differentiate individuals at a given point in development (see Forest et al., 2023; Saffran & Kirkham, 2018, for reviews of changes in regularity learning across development). Thus, we posit that substantial individual differences should be revealed in the ability to perceive and learn which continuously encountered variations in the input are relevant or informative and which are not.

This perspective also offers novel avenues for understanding the wide individual differences in response to interventions in clinical populations. ADHD (attention deficit hyperactivity disorder), for example, has been tied to heightened novelty seeking which leads to the suboptimal reward-related decision-making characteristic of this population (Lieder et al., 2019). For autism, findings have been mixed. Comparing trial-by-trial performance in a serial detection task, neurotypical participants were found to overweight recent statistics and quickly update their internal sensory models, which is adaptive in changing environments, whereas individuals with autism were found to rely atypically heavily on long-term statistics (Lieder et al., 2019). Other studies, however, provided evidence for faster rather than slower updating of internal models by individuals with autism (Goris et al., 2022; Lawson, Mathys, & Rees, 2017).

Role in the cognitive system

Here we come to an important distinction regarding cognitive architecture, and an important clarification: Per our Figure 1, IF does not preclude mechanisms of SL and does not aim to replace SL as a theoretical construct. Rather, while conventional SL theory regards SL as a stand-alone mechanism, from our perspective, SL computations form a subcomponent of IF, with SL playing a critical role in providing baselines against which change can be detected. This distinction is a corollary of the view of the environment as being fundamentally dynamic, with the *changes in regularities* as the main target

of perception. To identify such changes, a baseline of current covariation is needed, from which novel information can be detected. Hence, per our view, SL is a primary mechanism that non-selectively attunes the system to regularities in sensory input, but in service of the higher purpose of novelty detection. While we have argued so far that IF is a necessary component of learning regularities, as we outline below, from a wider perspective, we consider IF to be a unifying principle that underlies more complex behavior such as curiosity, exploration, and creativity. This perspective opens a novel set of questions regarding how SL and IF interact, and what outputs they produce, to enable efficient perception and action in a dynamic and everchanging environment.

One system or two?

An important theoretical question is whether registering recurrent regularities in the input for establishing baselines (SL), and sensitivity to changes in regularities that deviate from baseline (IF), require two independent and distinct systems, or whether one system can account for both. In other words, we ask whether a mechanism that registers recurrent regularities in sensory input as established by SL research could also at the same time allow for the fast accommodation of changes therein.

One could argue, for example, that a unified Bayesian perspective could, in principle, accommodate both SL and IF. In a Bayesian framework, if the environment is inherently everchanging, then the prior distribution of the perceived regularities would become less and less informative. With a relatively less informative distribution, deviating evidence (i.e., changes in regularity) has substantial weight, thereby significantly changing the posterior distribution and leading to increased sensitivity to novelty, aligning with our IF approach. In contrast, if the environment is stable and characterized by recurrent regularities, this would result in increasingly strong priors for expected events. Strong priors would lead to facilitation in the perception and/or processing of likely events in the input (e.g., Friston, 2005, 2009, 2010; Kok, Jehee, & de Lange, 2012), aligning with the conventional SL approach and its empirical findings. The main appeal of such an architecture is that one unified system generates different behaviors to the extent that the environment is stable versus ever-changing in a given domain.

While this Bayesian approach has the advantage of parsimony, it predicts, for example, that the more exposure a learner has to a repeated set of regularities, the less weight a surprising deviation from

this pattern will have in updating the posterior distribution. This is because repeated regularities produce increasingly strong priors, so that substantial evidence is eventually required for updating beliefs about the structural properties of the input.² This goes counter to the prediction we would make based on our IF account. IF argues that effective novelty detection and behavioral adaptation on the basis thereof, take place when there is a violation of structure in a so-far stable patterned environment (which has served as a baseline). In general, humans and other organisms adapt very rapidly to changing statistics while seeming to provisionally segregate changing statistics from regularities learned over the long term (e.g., Dell, Reed, Adams, & Meyer, 2000; Kraljic & Samuel, 2005; Onishi, Chambers & Fisher, 2002), and link them causally to contexts (e.g., Kraljic & Samuel, 2011). Evidence shows that detection and adaptation to deviation from recurrent regularities is exceedingly fast *regardless of lengthy past experience*. For example, participants were found to adjust their preferred saccade length (which reflected their prolonged reading experience) when presented with sentences containing words with a uniform word-length, when they had only one trial to adapt (Cutter et al., 2018).

While we opted to exemplify this problem using Bayesian terms, we should emphasize that virtually any learning system will face a challenge when it has been immersed in very strong regularities, and those regularities begin to change: Systems will respond sluggishly to changes in a previously highly regular environment. For example, a neural network (e.g., a simple recurrent network; Elman, 1990, 1991) may require substantial experience to overcome previous learning, and it might well lose significant aspects of prior learning if they are no longer reinforced (so-called *catastrophic interference*; McClosky & Cohen, 1989; Bower, Thomson-Schill, & Tulving, 1994; Merhav, Karni & Gilboa, 2014) — a phenomenon that is not generally observed in biology. This suggests that a system that targets the tracking of stable regularities, and a system that prioritizes novelty detection over stable regularities, might operate with different computational mechanisms.

This discussion resonates with our initial claim that a theory of learning regularities should first consider the nature of the environment that is the object of learning. We take it as evident that a

² It should be noted that the magnitude of surprise can be greater when encountering a highly unexpected input in the context of a strong prior (characterized by a narrow distribution) compared to a weak prior (characterized by a broad distribution). This is due to the potentially larger discrepancy between the unexpected input and the concentrated probability mass of a strong prior.

computational mechanism tracking regularities in the input could naturally assimilate gradual and slow changes, as original patterns will recur but at lower and lower frequency, and novel patterns will emerge at a higher and higher frequency. However, if the input is characterized by abrupt changes in regularities which would flag informational novelty, and the target of perception is precisely these changes, then slow adaptive learning mechanisms would probably not suffice.

This perspective raises important questions that should be the focus of future empirical and computational research efforts: What is the overlap in mechanisms that underlie the assimilation and updating of baseline regularities and those involved in detecting novel regularities? How do they interact with one another, and with what brain networks? Are different behavioral phenomena indicating sensitivity to novel regularities (spanning different cognitive domains and different timescales of adaptation) all tapping into one and the same IF system? Finally, are representations of baseline regularities overwritten/adjusted by IF, or do the representations of baseline regularities and those of novel regularities co-exist? As we outline in the following section, these questions regarding regularity learning are echoed across other domains in cognitive science.

Parallel debates across cognitive neuroscience

Since Helmholtz (1863), cognitive neuroscience has grappled with how sensory input and prior knowledge (expectations) interact to establish a percept (e.g., Geisler & Kersten, 2002; Friston, 2009; Heilbron, & Chait, 2018). Contemporary studies examine how expectations are established (Jabar & Fougne, 2022), how they shape behavior and neural response (Egger, Remington, Chang, & Jazayeri, 2019), and, importantly, how they change in a world with dynamic, fluctuating regularities (Hodson et al., 2023). Whereas this paper centers on SL, parallel debates on the importance of novelty are seen in multiple domains of cognition.

The most relevant framework is that of *predictive processing*. This theoretical approach to cognition views the brain as a Bayesian inference machine, where predictions regarding sensory input are continuously made to minimize free energy, a proxy for uncertainty and surprise, thereby facilitating perception and action (e.g., Friston, 2005, 2009, 2010). While we have discussed the problem of exceedingly fast adaptation with reference to a Bayesian perspective regarding enhanced sensitivity to

changes in regularities, parallels between current theories of predictive processing and our IF view certainly exist. Predictive processing, similar to IF, regards learners as active in the sense that they continuously make inferences, and it highlights the role of explorative behavior in learning (e.g., Friston, 2016, 2017; see Schwartenbeck et al., 2013, for discussion). Like IF, predictive processing centers on adaptation, offering computational mechanisms for it through the notion of minimizing prediction errors and a continuous process of updating priors. Given this, both frameworks reject the idea of an end state when the environmental regularities have been assimilated. However, an important basic difference remains. While models of predictive coding also offer computational accounts for novelty seeking, curiosity, and creativity (e.g., Schwartenbeck et al., 2013; see Clark, 2017, for discussion), in essence, they center on the notion of minimizing surprisal, while IF centers on the *prioritization* of changes in regularity. This distinction echoes our above discussion on one versus two systems.

In the domain of visual perception, the contrast between prioritizing expected input (as in the SL approach) vs. upweighting novel and surprising input (as in the proposed IF approach), is discussed in the context of sensory perception of unitary events vs. sensory-motor predictions. This debate, dubbed the *Perceptual Prediction Paradox* (see Press, Kok, & Yon, 2020), outlines the difference between, say, perceiving a cup, which requires fast recognition of a familiar object, vs. sensing the cup slipping from one's grip, which requires fast detection of deviation from what is expected regarding this object in terms of sensory information.

From a neurobiological perspective, a key signature of expectations is a weakened neural response to stimuli that are anticipated (see Heilbron & Chait, 2018; de Lange, Heilbron, & Kok, 2018, for reviews). There are, however, different accounts of the specific neural mechanism responsible for this suppression, which directly relate to whether perception tilts toward the input we expect, or if it instead prioritizes unexpected and novel input (see Press et al., 2020, for discussion). According to sharpening models, neural populations that are not tuned to the anticipated stimulus are particularly affected by expectations, resulting in a neural response that is overall diminished in magnitude but carries a more precise representation of the stimulus (e.g., Kok et al., 2012; Bell, Summerfield, Morin, Malecek, & Ungerleider, 2016). This “sharpening” process biases perception in accordance with the

perceiver's expectations, and is consistent with data demonstrating that predicted events are perceived with greater clarity and detected or processed faster, as typically assumed in SL theory. In contrast, damping models propose that neural populations that are tuned toward an expected stimulus are suppressed, leading to *reduced* cortical activation for expected input (Blakemore, Wolpert, & Frith, 1998; Summerfield & de Lange 2014). This leads to the prioritization of novelty, as the brain now favors the processing of surprising information, per an IF approach.

How these two seemingly incompatible explanations (and the data supporting each) can be reconciled is still a matter of debate (see, Richter, Heilbron, & de Lange, 2022, for discussion). While some theoretical accounts hypothesize that sharpening and damping can occur in parallel in different neural populations (Friston, 2005), others propose different time courses for sharpening and damping so that initial processing favors the expected, but later processing highlights signals that depart from these expectations to allow us to accommodate change (e.g., Press et al., 2020). These proposals for reconciliations, however, are silent regarding how the tradeoff between prediction and novelty detection is modulated by context, per our IF approach. The alignment of input statistics with active behavior may play an important role in such modulation (Roark, Lehet, Dick, & Holt, 2022; Lim, Fiez, & Holt, 2019).

Integrating SL into an IF framework sets a novel research agenda for understanding how input patterns acquired by SL may lead to predictions about upcoming input, and how departures from these patterns are flagged and accommodated in behavior and neural response. In the IF framework, future SL research will benefit from cross-fertilization with cognitive neuroscience literatures that examine the nature of predictive processes and novelty detection. Such research should focus on the precise neural systems that are implicated in dynamically tracking ever-changing regularities in sensory input, and to what extent they are related to reward networks driving curiosity and exploration.

Information foraging and higher cognitive functions

We consider the prioritization of novel information to be a domain-general feature and a unifying principle that explains a wide range of behaviors. Starting from early life, cognitive development is often cast in terms of constructing internal models that serve as a baseline for the

detection of important novel information (Atzil, Gao, Fradkin, & Barrett, 2018). As Twomey and Westermann (2018) suggest, infants drive their cognitive development by searching for structure in their environment, and maximal learning emerges when stimulus novelty is maximized in reference to their internal models. However, we propose that IF is also an explanatory principle for more complex behaviors. For example, foraging for information has been a cornerstone principle in the study of curiosity and its neurobiological underpinning (see e.g., Loewenstein, 1994; Kidd & Hayden, 2015, Gottlieb & Oudeyer, 2018, for review). Within this research area, novelty is taken to act as an intrinsic reward in exploration (e.g., Gottlieb, Oudeyer, Lopes, & Baranes, 2013), and such curiosity-driven explorations overall lead to improvement of prediction, a reduction in uncertainty, and assimilating more complex structures (Oudeyer & Smith, 2016). As long as three decades ago, tracking ocular movements, Berlyne (1996) found that when presented with pairs of stimuli, participants spend less and less time inspecting recurrent patterns and more and more time looking at novel patterns. While at the time, this was labeled “perceptual curiosity”, it coincides well with our IF approach. In general, current theories of curiosity converge on the assumption that the automatic bias towards novel and surprising events is rooted in the motivation to reduce uncertainty in the environment (van Lieshout, de Lange, & Cools, 2020), so that the model of the environment is continuously updated. This has been extensively shown in how young children forage for visual information. Infants tend to focus on familiar visual stimuli as long as they offer learning progress, but they switch to novel stimuli when learning progress drops (Poli et al., 2020; and see Adyman & Mareschal, 2013, for how redundancy governs spontaneous orientation). Indeed, if the environment is not stable but ever-changing, such updating is a primary priority. From this perspective, highly predictable events are uninformative and do not contribute to uncertainty reduction for updating our models of the world. Similar to our approach, novelty-based theories of curiosity suggest that new and highly uncertain stimuli drive curiosity, and, in general, the causal structure of the environment and its predictability will determine whether high or moderate uncertainty should drive curiosity (Dubey & Griffiths, 2020).

Discussions in the domain of creativity parallel the proposed contrast between SL and IF. In creativity research, “creative foraging”, as a theoretical construct, is taken to balance two main processes, exploitation and exploration (Hills, Todd, Lazer, Redish, & Couzin, 2015; Hart et al., 2018).

Exploitation is operationally defined as taking advantage of the specific regularities within a search space, repeatedly applying identical or similar computations. Opposite to exploitation is exploration, defined as moving to a novel search space, applying novel computations, to increase gain. While exploitation would maximize reward if the environment is stable, a maximizing organism would not detect superior reward associated with different regions of space or different computations (indeed, exploitation may limit reward, if resources are depleted, especially if multiple organisms compete for resources; e.g., Gallistel, 1993). While exploitation is driven by predictable outcomes akin to SL, exploration, and in essence, curiosity behavior, is driven by the promise of information gain that would come in unpredictable novel search spaces (Linquin & Lombrozo, 2020). In fact, exploration is defined only with reference to exploitation, just as novelty in IF is defined with reference to baseline statistical regularity.

Concluding remarks and future directions

Our present theoretical perspective on IF and its role in cognitive architecture lays the groundwork and raises new questions for future research on the processing of regularities. Critically, experimental investigations of IF should focus on paradigms that involve *changes* in regularity, tracking participants' perception of these changes in real-time, mapping the precise parameters that determine efficient detection of deviations from baseline (e.g., Weiss, Gerfen, & Mitchel, 2009; Ryskin, Qi, Duff, & Brown-Schmidt, 2017; Wang & Theeuwes, 2020; Hodson et al., 2023). When it comes to laboratory experiments with artificial stimuli, this will require tracking behavior in significantly longer experimental sessions than conventional SL research (see e.g., Frank, Tannenbaum, & Gibson, 2013), where participants are processing input streams that vary in regularity as the session proceeds. In the domain of language, where the statistical co-occurrences of linguistic elements can be determined by considering large databases, experimental work can focus on presenting participants with input streams that conform or not with the distributional properties that characterize their linguistic environment. This approach can be used to measure the manipulation's impact on performance, for speech, print, or any linguistic input (e.g., Idemaru & Holt, 2011; Isbilen, McCauley, & Christiansen, 2022; and see Elazar, Alhama, Bogaerts, & Frost, 2022, for discussion). In the same vein, from the perspective of individual

differences, studies should focus on individual sensitivity to a change in statistical regularities, individual plasticity in adapting to novel structural properties, and individual efficiency in determining which regularities are informative and which are not, given the particular context of learning.

To map the neurobiological underpinning of IF, research could focus on the range of neural mechanisms that are tuned to track deviation from baseline, and mechanisms that flag alterations in patterns of quasi-regularity. A recent example is the role of norepinephrine in tracking unexpected uncertainty and deviation from regularity (Zhao et al., 2019). Importantly, such investigation should go beyond simple oddball paradigms and their concurrent mismatch negativity responses (e.g., Näätänen & Alho, 1995). As we point out above, and in contrast to conventional SL, mechanisms of IF most probably involve complex interactions with systems that govern exploration, attention, and reward. Establishing the neurocircuitry that is implicated in the foraging of change, and the necessary conditions for its consolidation in memory, is an important priority.

References

- Addyman, C. & Mareschal, D. (2013) Local redundancy governs infants' spontaneous orienting to visual-temporal sequences *Child Development*, *84*, 1137-1144.
- Arnal, L. H., & Giraud, A. L. (2012). Cortical oscillations and sensory predictions. *Trends in cognitive sciences*, *16*, 390-398.
- Asr, F. T., & Demberg, V. (2015, April). Uniform surprisal at the level of discourse relations: Negation markers and discourse connective omission. *Proceedings of the 11th international conference on computational semantics*. Pp. 118-128.
- Atzil, S., Gao, W., Fradkin, I., & Barrett, L. F. (2018). Growing a social brain. *Nature Human Behaviour*, *2*, 624-636.
- Aylett, M., & Turk, A. (2004). The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech. *Language and speech*, *47*, 31-56.
- Behrens TE, Woolrich MW, Walton ME, & Rushworth MF. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*, 1214-1221.
- Bell, A. H., Summerfield, C., Morin, E. L., Malecek, N. J., & Ungerleider, L. G. (2016). Encoding of Stimulus Probability in Macaque Inferior Temporal Cortex. *Current Biology*, *26*, 2280–2290.
- Berlyne, D. E. (1966). Curiosity and Exploration: Animals spend much of their time seeking stimuli whose significance raises problems for psychology. *Science*, *153*(3731), 25-33.
- Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature neuroscience*, *1*, 635–640.
- Bock, J. K. (1986). Syntactic persistence in language production. *Cognitive psychology*, *18*, 355-387.
- Bogaerts, L., Richter, C. G., Landau, A. N., & Frost, R. (2020). Beta-Band Activity Is a Signature of Statistical Learning. *The Journal of neuroscience*, *40*, 7523–7530.
- Bogaerts, L., van Moorselaar, D., & Theeuwes, J. (2022). Does it help to expect distraction? Attentional capture is attenuated by high distractor frequency but not by trial-to-trial

predictability. *Journal of Experimental Psychology: Human Perception and Performance*, 48, 246–261.

Bower, G. H., Thompson-Schill, S., & Tulving, E. (1994). Reducing retroactive interference: An interference analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(1), 51–66.

Braunlich, K., & Love, B. C. (2022). Bidirectional influences of information sampling and concept learning. *Psychological Review*, 129, 213–234.

Bröker, F., Holt, L. L., Roads, B. D., Dayan, P., & Love, B. C. (2024). Demystifying unsupervised learning: how it helps and hurts. *Trends in Cognitive Sciences*.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.

Brown-Schmidt, S., & Tanenhaus, M. K. (2008). Real-time investigation of referential domains in unscripted conversation: A targeted language game approach. *Cognitive science*, 32, 643-684.

Bruner, J. S., & Postman, L. (1949). On the perception of incongruity: A paradigm. *Journal of Personality*, 18(2), 206-223.

Ryskin, R. A., Yoon, S. O., & Brown-Schmidt, S. (2015). Language, perspective, and memory. In R. A. Scott & S. M. Kosslyn, (2015). *Emerging trends in the social and behavioral sciences: An interdisciplinary, searchable, and linkable resource*. John Wiley & Sons.

Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition*, 121, 127-132.

Christiansen, M.H. & Chater, N. (2016). The Now-or-Never bottleneck: A fundamental constraint on language. *Behavioral & Brain Sciences*, 39, e62.

Christiansen, M.H. & Chater, N. (2022). *The language game: How improvisation created language and changed the world*. Basic Books.

Contreras Kallens, P., Kristensen-McLachlan, R. D., & Christiansen, M. H. (2023). Large language models demonstrate the potential of statistical learning in language. *Cognitive Science*, 47(3), e13256.

Conway, C.M. and M.H. Christiansen. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 24-39.

Cutter, M.G., Drieghe, D., & Liversedge, S.P. (2017). Reading sentences of uniform word length: evidence for the adaptation of the preferred saccade length during reading. *Journal of Experimental Psychology: Human Perception and Performance*, 43, 1895- 1911.

Cutter, M.G., Drieghe, D., & Liversedge, S.P. (2018). Reading sentences of uniform word length – II: Very rapid adaptation of the preferred saccade length. *Psychonomic Bulletin & Review*, 5, 1435-1440.

De Lange, F. P., Heilbron, M., & Kok, P. (2018). How Do Expectations Shape Perception? *Trends in cognitive sciences*, 22(9), 764–779.

Dell, G. S., Reed, K. D., Adams, D. R., & Meyer, A. S. (2000). Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1355–1367.

Dideriksen, C., Christiansen, M.H., Tylén, K., Dingemanse, D. & Fusaroli, R. (2023). Quantifying the interplay of conversational devices in building mutual understanding. *Journal of Experimental Psychology: General*, 152, 864–889.

Dubey, R., & Griffiths, T. L. (2020). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, 127, 455-476.

Egger, S. W., Remington, E. D., Chang, C. J., & Jazayeri, M. (2019). Internal models of sensorimotor integration regulate cortical dynamics. *Nature neuroscience*, 22, 1871–1882.

Egner, T., Monti, J. M., & Summerfield, C. (2010). Expectation and surprise determine neural population responses in the ventral visual stream. *Journal of Neuroscience*, 30(49), 16601-16608.

Eimas, P.D., & Corbit, J.D. (1973). Selective adaptation of linguistic feature detectors. *Cognitive Psychology*, 4, 99-109.

- Eisner, F., McQueen, J.M. (2005). The specificity of perceptual learning in speech processing. *Perception & Psychophysics* 67, 224–238
- Elazar, A., Alhama, R., Bogaerts, L., et al. (2022). When the “Tabula” is Anything but “Rasa:” What Determines Performance in the Auditory Statistical Learning Task? *Cognitive Science*, 46.
- Elman, J.L. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- Elman, J.L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7, 195-224.
- Emberson, L.L., Liu, R., & Zevin J.D. (2013). Is statistical learning constrained by lower level perceptual organization? *Cognition*, 128, 82-102.
- Ferguson, B., & Lew-Williams, C. (2016). Communicative signals support abstract rule learning by 7-month-old infants. *Scientific Reports*, 6, doi: 10.1038/srep25434.
- Fiser, J. and R.N. Aslin (2001) *Unsupervised statistical learning of higher-order spatial structures from visual scenes. Psychological Science*, 12, 499-504.
- Fiser, J. and R.N. Aslin. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 458-467.
- Forest, T. A., Siegelman, N., & Finn, A. S. (2022). Attention Shifts to More Complex Structures with Experience. *Psychological Science*, 33(12), 2059–2072.
- Forest, T.A., Schlichting, M.L., Duncan, K., & Finn, A.S. (2023). How statistical learning changes across development. *Nature Reviews Psychology*, 2, 205-219.
- Fowler, C. A., & Housum, J. (1987). Talkers’ signaling of “new” and “old” words in speech and listeners’ perception and use of the distinction. *Journal of memory and language*, 26, 489-504.
- Frank, A. F., & Jaeger, T. F. (2008). Speaking rationally: Uniform information density as an optimal strategy for language production. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 30, No. 30).
- Frank M.C., Tenenbaum, J.B., & Gibson, E. (2013). Learning and Long-Term Retention of Large-Scale Artificial Languages. *PLOS ONE* 8(4): 10.1371/annotation/fcf19f94-5f2c-4a5f-9a6f-3d880bcdd850.

- Friston, K.J. (2005). A theory of cortical responses. *Philosophical Transaction of the Royal Society: Biology*, 360, 815-836
- Friston K. (2009). The free-energy principle: a rough guide to the brain?. *Trends in Cognitive Sciences*, 13, 293–301.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2), 127-138.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., O’Doherty, J. and Pezzulo, G. (2016) Active Inference and Learning. *Neuroscience & Biobehavioral Reviews*, 68, 862-879.
- Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, Pezzulo G. (2017). Active Inference: A Process Theory. *Neural Computation*, 29, 1-51.
- Frost, R., Armstrong, B.C., & Christiansen, M.H. (2019). Statistical learning research: A critical review and possible new directions. *Psychological Bulletin*, 145, 1128–1153.
- Frost, R., Armstrong, B.C., Siegelman, N., & Christiansen, M.H. (2015). Domain generality vs. modality specificity: The paradox of statistical learning. *Trends in Cognitive Science*, 19, 117-125.
- Fusaroli, R., Rączaszek-Leonardi, J., & Tylén, K. (2014). Dialog as interpersonal synergy. *New Ideas in Psychology*, 32, 147–157.
- Gabay, Y., Dick, F. K., Zevin, J. D., & Holt, L. L. (2015). Incidental auditory category learning. *Journal of experimental psychology. Human perception and performance*, 41, 1124–1138.
- Gabay, Y., Karni, A., & Holt, L. L. (2023). Memory for incidentally learned categories evolves in the post-learning interval. *eLife*, 12, e81855.
- Gallistel, R. C. (1993). *The Organization of Learning*. Cambridge, MA: MIT Press.
- Garrod, S., & Pickering, M. J. (2004). Why is conversation so easy? *Trends in cognitive sciences*, 8, 8-11.
- Gati, I., & Ben-Shakhar, G. (1990). Novelty and significance in orientation and habituation: a feature-matching approach. *Journal of Experimental Psychology: General*, 119, 251.
- Geisler W. S., Kersten D. (2002). Illusions, perception and Bayes. *Nature Neuroscience*. 5, 508–510.

- Geng, J. J., Won, B. Y., & Carlisle, N. (2019). Distractor ignoring: Strategies, learning, and passive filtering. *Current Directions in Psychological Science*, 28, 600–606.
- Genzel, D., & Charniak, E. (2002). Entropy rate constancy in text. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics* (pp. 199-206).
- Gibson, E., Futrell, R., Piantadosi, S.P. Dautriche, I., Mahowald, K., Bergen, L., Levy, R. (2019). How Efficiency Shapes Human Language. *Trends in Cognitive Sciences*, 23, 389-407.
- Goris, J., Braem, S., Van Herck, S., Simoens, J., Deschrijver, E., Wiersema, J.R., Paton, B., Brass, M., Todd, J. (2022). Reduced Primacy Bias in Autism during Early Sensory Processing. *Journal of Neuroscience*, 42, 3989-3999.
- Gottlieb, G. (2007). Probabilistic epigenesis. *Developmental science*, 10, 1-11.
- Gottlieb, J., Oudeyer, P. Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in cognitive sciences*, 17, 585-593.
- Gottlieb, J., & Oudeyer, P. Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*, 19, 758-770.
- Goulart, L., Gray, B., Staples, S., Black, A., Shelton, A., Biber, D., Egbert, J. & Wizner, S. (2020). Linguistic perspectives on register. *Annual Review of Linguistics*, 6, 435–455.
- Gruber, M.J., & Ranganath, C. (2019). How Curiosity Enhances Hippocampus-Dependent Memory: The Prediction, Appraisal, Curiosity, and Exploration (PACE) Framework. *Trends in Cognitive Sciences*, 23, 1014-1025.
- Hart, Y., Goldberg, H., Striem-Amit, E., Mayo, A. E., Noy, L., & Alon, U. (2018). Creative exploration as a scale-invariant search on a meaning landscape. *Nature communications*, 9(1), 5411.
- Hasson, U. (2017). The neurobiology of uncertainty: implications for statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160048.
- Heilbron, M., & Chait, M. (2018). Great Expectations: Is there Evidence for Predictive Coding in Auditory Cortex? *Neuroscience*, 389, 54–73.
- Helmholtz, H. von (1863). Die Lehre von den Tonempfindungen als physiologische Grundlage für die Theorie der Musik (The Study of the Sensations of Tone as a Physiological Foundation for Music Theory) (in German) (1 ed.).

- Henin, S., Turk-Browne, N. B., Friedman, D., Liu, A., Dugan, P., Flinker, A., ... & Melloni, L. (2021). Learning hierarchical sequence representations across human cortex and hippocampus. *Science advances*, *7*, eabc4530.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in cognitive sciences*, *19*, 46-54.
- Hipp, J. F., Engel, A. K., & Siegel, M. (2011). Oscillatory synchronization in large-scale cortical networks predicts perception. *Neuron*, *69*, 387-396.
- Hodson, A.J., Shinn-Cunningham, B.G., & Holt, L. L. (2023). Statistical learning across passive listening adjusts perceptual weights of speech. PsyArxiv, <https://psyarxiv.com/4kxz3/>.
- Holt, L. L. (2005). Temporally non-adjacent non-linguistic sounds affect speech categorization. *Psychological Science*, *16*, 305-312.
- Huang, J. & Holt, L.L. (2012). Listening for the norm: adaptive coding in speech categorization. *Frontiers in Perception Science*, *3:10*. PMC3078024.
- Idemaru, K. & Holt, L. L. (2011). Word recognition reflects dimension-based statistical learning. *Journal of Experimental Psychology: Human Perception & Performance*, *37*, 1939-1956. PMC3285244.
- Isbilen, E.S., McCauley, S.M. & Christiansen, M.H. (2022). Individual differences in artificial and natural language statistical learning. *Cognition*, *225*, 105123.
- Jabar, S. B., & Fougny, D. (2022). How do expectations change behavior? Investigating the contributions at encoding versus decision-making. *Journal of experimental psychology. Learning, Memory, and Cognition*, *48*, 226.
- Johnsrude, I. S., Mackey, A., Hakyemez, H., Alexander, E., Trang, H. P., & Carlyon, R. P. (2013). Swinging at a Cocktail Party: Voice familiarity aids speech perception in the presence of a competing voice. *Psychological Science*, *24*, 1995-2004.
- Karmazyn-Raz, H., & Smith, L. B. (2023). Sampling statistics are like story creation: a network analysis of parent–toddler exploratory play. *Philosophical Transactions of the Royal Society B*, *378*, 20210358.

- Karuza, E. A., Li, P., Weiss, D. J., Bulgarelli, F., Zinszer, B. D., & Aslin, R. N. (2016). Sampling over nonuniform distributions: A neural efficiency account of the primacy effect in statistical learning. *Journal of Cognitive Neuroscience*, *28*(10), 1484–1500.
- Kemény, F., & Németh, K. (2017). Stimulus dependence and cross-modal interference in sequence learning. *The Quarterly Journal of Experimental Psychology*, *70*, 2535–2547.
- Kidd, C., & Hayden, B.Y. (2015). The psychology and neuroscience of curiosity, *Neuron*, *88*, 449-460.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS ONE*, *7*(5), e36399.
- Kim, M., Horton, W., and Bradlow, A. R. (2011). Phonetic accommodation between native and non-native speakers. *Journal of Laboratory Phonology*, *2*, 125–156.
- Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & memory*, *1*, 106-120.
- Kok, P., Jehee, J.F.M, & de Lange, F.P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, *75*, 265-270.
- Kolling, N., Wittmann, M., Behrens, T. et al. (2016). Value, search, persistence and model updating in anterior cingulate cortex. *Nature Neuroscience*, *19*, 1280–1285.
- Kraljic, T., & Samuel, A. G. (2005). Perceptual learning for speech: is there a return to normal? *Cognition*, *51*, 141-178.
- Kraljic, T., & Samuel, A. G. (2011). Perceptual learning evidence for contextually-specific representations. *Cognition*, *121*, 459-465.
- Kraljic, T., Samuel, A. G., & Brennan, S. E. (2008). First impressions and last resorts: How listeners adjust to speaker variability. *Psychological Science*, *19*, 332-338.
- Krauss, R. M., Freyberg, R., & Morsella, E. (2002). Inferring speakers' physical attributes from their voices. *Journal of Experimental Social Psychology*, *38*, 618-625.

Kumar, S., Kaposvari, P., & Vogels, R. (2017). Encoding of predictable and unpredictable stimuli by inferior temporal cortical neurons. *Journal of Cognitive Neuroscience*, 29(8), 1445-1454.

Lawson, R., Mathys, C., & Rees, G. (2017). Adults with autism overestimate the volatility of the sensory environment. *Nature Neuroscience*, 20, 1293–1299.

Lee, Y., Gordon Danner, S., Parrell, B., Lee, S., Goldstein, L., & Byrd, D. (2018). Articulatory, acoustic, and prosodic accommodation in a cooperative maze navigation task. *Plos one*, 13(8), e0201444.

Leech, R., Holt, L. L., Devlin, J. T., & Dick, F. (2009). Expertise with artificial nonspeech sounds recruits speech-sensitive cortical regions. *The Journal of neuroscience*, 29, 5234–5239.

Lelonkiewicz, J. R., Ktori, M., & Crepaldi, D. (2020). Morphemes as letter chunks: Discovering affixes through visual regularities. *Journal of Memory and language*, 115, 104-152.

Li, A.S., van Moorselaar, D. ,& Theeuwes, J.(2024). Attending is not enough: Responding to targets is needed for across-trial statistical learning. *Attention, Perception, & Psychophysics*. doi.org/10.3758/s13414-024-02952-0

Lieder, I., Adam, V., Frenkel, O., Jaffe-Dax, S., Sahani, M., & Ahissar, M. (2019). Perceptual bias reveals slow updating in autism and fast-forgetting in dyslexia. *Nature Neuroscience*, 22, 256–264.

Lim, S.J., Fiez, J.A., Holt, L.L. (2019). Role of the striatum in incidental learning of sound categories. *Proceedings of the National Academy of Sciences*, 116, 4671-4680.

Linquin, E.g., & Lombrozo, T. (2020). Explanation-seeking curiosity in childhood, *Current Opinion in Behavioral Sciences*, 35, 14–20.

Loewenstein, G. (1994). The Psychology of Curiosity: A Review and Reinterpretation, *Psychological Bulletin*, 116, 75-98.

Lu, K., & Vicario, D. S. (2014). Statistical learning of recurring sound patterns encodes auditory objects in songbird forebrain. *Proceedings of the national academy of sciences*, 111, 14553-14558.

Luthra, S. (2023). Why are listeners hindered by talker variability?. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-023-02355-6>

Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, *283*, 77-80.

Marcus, G. F., Fernandes, K. J., & Johnson, S. P. (2007). Infant rule learning facilitated by speech. *Psychological Science*, *18*, 387–391.

Mareschal, D., & French, R. M. (2017). TRACX2: a connectionist autoencoder using graded chunks to model infant visual statistical learning. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *372*(1711), 20160057.

McCloskey, M., & Cohen, N. J. (1989). Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 24, pp. 109-165). Academic Press.

Meyer, T., & Olson, C. R. (2011). Statistical learning of visual transitions in monkey inferotemporal cortex. *Proceedings of the National Academy of Sciences*, *108*(48), 19401-19406.

Misyak, J. B., Christiansen, M. H., & Tomblin, J. B. (2010). On-line individual differences in statistical learning predict language processing. *Frontiers in Psychology*, *1*, 31.

Munson, B. & Babel, M. (2007), Loose lips and silver tongues, or, projecting sexual orientation through speech. *Language and Linguistics Compass*, *1*, 416-449. <https://doi.org/10.1111/j.1749-818X.2007.00028.x>

Näätänen, R., & Alho, K. (1995). Mismatch negativity: A unique measure of sensory processing in audition. *International Journal of Neuroscience*, *80*, 317-337.

Nastase, S., Iacovella, V., & Hasson, U. (2014). Uncertainty in visual and auditory series is coded by modality-general and modality-specific neural systems. *Human Brain Mapping*, *35*, 1111–1128.

Norris, D., McQueen, J.M., & Cutler, A. (2003). Perceptual learning in speech. *Cognitive Psychology*, *47*, 204-238.

Nygaard, L. C., & Pisoni, D. B. (1998). Talker-specific learning in speech perception. *Perception & psychophysics*, *60*(3), 355-376.

Onishi, K. H., Chambers, K. E., & Fisher, C. (2002). Learning phonotactic constraints from brief auditory experience. *Cognition*, 83, B13-B23.

Oudeyer, P.Y., & Smith, L.B. (2016). How evolution may work through curiosity-driven developmental process. *Topics in Cognitive Science*, 8, 492-502.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). *Training language models to follow instructions with human feedback* (arXiv:2203.02155). arXiv. <https://doi.org/10.48550/arXiv.2203.02155>.

Pardo, J. S. (2006). On phonetic convergence during conversational interaction. *The Journal of the Acoustical Society of America*, 119, 2382-2393.

Pavlov, I. P. (1927). *Conditioned reflexes: an investigation of the physiological activity of the cerebral cortex*. Oxford Univ. Press.

Perrachione, T. K., Pierrehumbert, J. B., & Wong, P. (2009). Differential neural contributions to native-and foreign-language talker identification. *Journal of Experimental Psychology: Human Perception and Performance*, 35(6), 1950.

Pirulli, P., & Card, S. (1999). Information foraging. *Psychological Review*, 106, 643–675.

Press, C, Kok, P, &Yon, D. (2020). The perceptual prediction paradox. *Trends in Cognitive Science*.

Poli, F., Serino, G., Mars, R.B., & Hunnius, S. (2020). Infants tailor their attention to maximize learning. *Science Advances*, 6, DOI: [10.1126/sciadv.abb5053](https://doi.org/10.1126/sciadv.abb5053).

Rastle, K., Lally, C., Davis, M. H., & Taylor, J. S. H. (2021). The dramatic impact of explicit instruction on learning to read in a new writing system. *Psychological Science*, 32(4), 471-484.

Richter, D., Ekman, M., & de Lange, F.P. (2018). Suppressed Sensory Response to Predictable Object Stimuli throughout the Ventral Visual Stream. *Journal of Neuroscience*, 38, 7452-7461.

Richter, D., Heilbron, M., de Lange, F. P. (2022). Dampened sensory representations for expected input across the ventral visual stream, *Oxford Open Neuroscience*, 1, kvac013.

Roark, C.L., Lehet, M.I., Dick, F., Holt, L.L. (2022). The representational glue for incidental category learning is alignment with task-relevant behavior. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 48, 769-784.

Ryskin, R. A., Qi, Z., Duff, M. C., & Brown-Schmidt, S. (2017). Verb biases are shaped through lifelong learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(5), 781–794. <https://doi.org/10.1037/xlm0000341>

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

Saffran, J. R., & Kirkham, N. Z. (2018). Infant statistical learning. *Annual Review of Psychology*, 69, 181–203.

Saffran, J. R., Newport, E. L., Aslin, R. N., Tunick, R. A., & Barrueco, S. (1997). Incidental language learning: Listening (and learning) out of the corner of your ear. *Psychological Science*, 8, 101-105.

Saffran, J. R., Pollak, S. D., Seibel, R. L., & Shkolnik, A. (2007). Dog is a dog is a dog: Infant rule learning is not specific to language. *Cognition*, 105(3), 669-680.

Samuel, A.G. (1986). Red herring detectors and speech perception: In defense of selective adaptation. *Cognitive Psychology*, 18, 452-499.

Santolin, C., & Saffran, J. R. (2018). Constraints on Statistical Learning Across Species. *Trends in Cognitive Sciences*, 2, 52–63.

Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160049.

Schapiro, A., & Turk-Browne, N. (2015). Statistical Learning. *Brain Mapping: An Encyclopedic Reference*, 3, 501–506.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593-1599.

Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., ... & Greicius, M. D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. *Journal of Neuroscience*, *27*, 2349-2356.

Sethi, A., Voon, V., Critchley, H. D., Cercignani, M., & Harrison, N. A. (2018). A neurocomputational account of reward and novelty processing and effects of psychostimulants in attention deficit hyperactivity disorder. *Brain*, *141*, 1545-1557.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal*, *27*, 379-423.

Sheya, A., & Smith, L. (2019). Development weaves brains, bodies and environments into cognition. *Language, cognition and neuroscience*, *34*, 1266-1273.

Schwartenbeck, P., FitzGerald, T., Dolan, R. J., & Friston, K. (2013). Exploration, novelty, surprise, and free energy minimization. *Frontiers in psychology*, *4*, 710.

Siegelman, N., & Frost, R. (2015). Statistical learning as an individual ability: Theoretical perspectives and empirical evidence. *Journal of memory and language*, *81*, 105–120.

Siegelman, N., Bogaerts, L., Armstrong, B. C., & Frost, R. (2019). What exactly is learned in visual statistical learning? Insights from Bayesian modeling. *Cognition*, *192*, 104002.

Siegelman, N., Bogaerts, L., Christiansen, M. H., & Frost, R. (2017). Towards a theory of individual differences in statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1711), 20160059.

Siegelman, N., Bogaerts, L., Kronenfeld, O., & Frost, R. (2018). Redefining "Learning" in Statistical Learning: What Does an online measure reveal about the assimilation of visual regularities? *Cognitive science*, *42*, 692–727.

Siegelman, N., Rueckl, J. G., Steacy, L. M., Frost, S. J., van den Bunt, M., Zevin, J. D., ... & Morris, R. D. (2020). Individual differences in learning the regularities between orthography, phonology and semantics predict early reading skills. *Journal of memory and language*, *114*, 104145.

Siegelman, N., Kearns, D. M., & Rueckl, J. G. (2020). Using information-theoretic measures to characterize the structure of the writing system: The case of orthographic-phonological regularities in English. *Behavior Research Methods*, *52*, 1292–1312.

Sirois, S., & Mareschal, D. (2004). An interacting systems model of infant habituation. *Journal of cognitive neuroscience*, *16*, 1352-1362.

Slone, L. K., Abney, D. H., Smith, L. B., & Yu, C. (2023). The temporal structure of parent talk to toddlers about objects. *Cognition*, *230*, 105266.

Smith, L. B., Jayaraman, S., Clerkin, E. & Yu, C. (2018). The developing infant creates a curriculum for statistical learning. *Trends in Cognitive Sciences*, *4*, 325-336.

Smith, L. B. & Thelen, E. (2003) Development as a dynamic system. *Trends in Cognitive Science*, *7*, 343-348.

Snell, J., & Theeuwes, J. (2020). A story about statistical learning in a story: Regularities impact eye movements during book reading. *Journal of Memory and Language*, *113*, 104-127.

Sokolov, E. N. (1963). Higher nervous functions: The orienting reflex. *Annual review of physiology*, *25*, 545-580.

Suanda, S. H., Barnhart, M., Smith, L. B., & Yu, C. (2019). The signal in the noise: The visual ecology of parents' object naming. *Infancy*, *24*, 455-476.

Suanda, S. H., Smith, L. B. & Yu, C. (2017). The multisensory mature of verbal discourse in parent-toddler interactions. *Developmental Neuropsychology*, *41*, 324-341.

Suarez-Rivera, C., Smith, L. B. & Yu, C. (2019). Multimodal parent behaviors within joint attention support sustained attention in infants. *Developmental Psychology*, *55*, 96-109.

Summerfield C. & de Lange F.P. (2014) Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*:745–756.

Tandoc, M. C., Nadendla, B., Pham, T., & Finn, A. S. (2024). Directing Attention Shapes Learning in Adults but Not Children. *Psychological Science*. doi/10.1177/09567976241263347.

Theeuwes, J., Bogaerts, L., & van Moorselaar, D. (2022). What to expect where and when: How statistical learning drives visual selection. *Trends in Cognitive Sciences*, *26*(10), 860–872.

Trujillo, J.P., Dideriksen, C., Tylén, K., Christiansen, M.H. & Fusaroli, R. (2023). The dynamic interplay of kinetic and linguistic coordination in Danish and Norwegian conversation. *Cognitive Science*, 47, e13298.

Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The Automaticity of Visual Statistical Learning. *Journal of Experimental Psychology: General*, 134, 552-564.

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

Twomey, K.E. & Westermann, G. (2018). Curiosity-based learning in infants: a neurocomputational approach. *Developmental Science*, 21, e12629 doi: [10.1111/desc.12629](https://doi.org/10.1111/desc.12629).

van Lieshout, L. L., de Lange, F. P., & Cools, R. (2020). Why so curious? Quantifying mechanisms of information seeking. *Current Opinion in Behavioral Sciences*, 35, 112-117.

Vassena, E., Holroyd, C.B. and Alexander, W.H., 2017. Computational models of anterior cingulate cortex: At the crossroads between prediction and effort. *Frontiers in neuroscience*, 11, p.316.

Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychonomic Bulletin & Review*, 19(5), 871-878.

Wang, B., & Theeuwes, J. (2020). Implicit attentional biases in a changing environment. *Acta psychologica*, 206, 103064.

Weiss, D. J., Gerfen, C., & Mitchel, A. D. (2009). Speech segmentation in a simulated bilingual environment: A challenge for statistical learning? *Language learning and development*, 5, 30–49.

Wells, J., Christiansen, M.H., Race, D.S., Acheson, D. & MacDonald, M.C. (2009). Experience and sentence processing: Statistical learning and relative clause comprehension. *Cognitive Psychology*, 58, 250-271.

Werker, J. F., Gilbert, J. H., Humphrey, K., & Tees, R. C. (1981). Developmental aspects of cross-language speech perception. *Child development*, 349-355.

Won, B. Y., Kosoyan, M., & Geng, J. J. (2019). Evidence for second-order singleton suppression based on probabilistic expectations. *Journal of Experimental Psychology: Human Perception and Performance*, 45(1), 125–138.

Yan, C., Ehinger, B. V, Pérez-Bellido, A., Peelen, M. V, & de Lange, F. P. (2023). Humans predict the forest, not the trees: statistical learning of spatiotemporal structure in visual scenes. *Cerebral Cortex*, bhad115.

Yan, S., & Jaeger, F.T. (2020) Expectation adaptation during natural reading. *Language, Cognition and Neuroscience*, 35, 1394-1422.

Yu, C., & Smith, L. B. (2007). Rapid word learning under uncertainty via cross-situational statistics. *Psychological science*, 18, 414–420.

Yu, C., Zhang, Y., Slone, L. K., & Smith, L. B. (2021). The infant's view redefines the problem of referential uncertainty in early word learning. *Proceedings of the National Academy of Sciences*, 118, e2107019118.

Zhao, S. Chait, M., Dick, F., Dayan, P. Furukawa, S, Liao, H-I, (2019). Pupil-linked phasic arousal evoked by violation but not emergence of regularity within rapid sound sequences. *Nature Communication*, 10:4030.