The Acquisition Process

To appear in
The Cambridge Companion to Child Syntax

Lisa Pearl
University of California, Irvine

Abstract

This chapter discusses the acquisition process, and how we build theories that explain how that marvelous process occurs in children. I review insights about the relevant factors in the acquisition process, and what we currently know about those factors. I draw on results from quantitative approaches to language acquisition that allow us to concretely articulate and evaluate potential acquisition theories, such as mathematical learnability and computational modeling. The insights gained from these approaches will hopefully allow us to construct a complete, elegant theory of syntactic language acquisition.

1 Explaining the marvelous process of language acquisition

As scientists who care about understanding language acquisition, we'd like to identify key concepts that we need in order to create insightful, elegant theories of language acquisition. First of all, why exactly do we care about language acquisition? And second, why do we want to create a theory about it? Let's start with our interest in acquisition.

Put simply, language acquisition is something of a marvel that seems in need of explanation. The knowledge that children develop about their native language(s) is quite sophisticated, and children develop it more completely than adults learning a new language typically can. The time children take to develop this sophisticated knowledge is relatively short, and in fact shorter than the time taken by many adults who try to develop proficiency in a non-native language. The cognitive capacities (like attention, memory, executive control, and so on) that children have available to deploy towards the development of this sophisticated knowledge are also limited, especially when compared to adult capacities.

Moreover, children are often doing other things at the same time as they're developing their linguistic knowledge: teasing a sibling, claiming a coveted space on a parent's lap, or learning to barter for a desired toy (among many other activities). Children aren't sitting down with laser focus to analyze their input for relevant linguistic information the way adults might do in structured educational settings. Yet, children achieve native proficiency while most adults learning a new language don't. In other words, children succeed at language acquisition while most adults don't. This is a marvel. Importantly, the considerations I mentioned above (the nature of children's knowledge; their limited time, cognitive capacities, and focus) are only some of the factors that we need to keep in mind when coming up with an explanation for how this marvelous process occurs.

So, how *do* we come up with an explanation? A theory is an explanation, and so that's why we want a theory of acquisition. With this goal in mind, how do we come up with a theory about how the process of acquisition works? If we want to build a theory, we need to be precise about what we're trying to explain and how we're going to try to explain it. More specifically, for language acquisition, we're trying to explain the empirical data available about how language acquisition works (what children seem to have acquired when, and how they seem to do it). A theory is a useful abstraction of those empirical data (Dennett, 1990, Futrell & Mahowald, 2025) – that is, a theory is a way to "compress" the data so we can explain how those data were generated in a way that doesn't involve "cumbersomeness, lack of generality, and unwanted detail" (Dennett, 1990).

To be clear, it's not that we want to ignore empirical data we view as valid. Rather, we want to find a way to capture those data in some sort of "compressed" way. Let me illustrate with a toy example. Suppose we want to capture the following data about the target linguistic knowledge for a child:

(1) *wh*-dependency knowledge
- a. Allowed
  - (i) Who did she think the kitty was for $\_\_{who}$?
  - (ii) What did she think $\_\_{what}$ was pretty?
- b. Disallowed
  - (i) Who did she think the kitty for $\_\_{who}$ was pretty?

An uncompressed "explanation" like (2a) might just recapitulate the data directly. This "explanation" seems fairly cumbersome with a lot of extra detail, and this explanation won't generalize to new *wh*-dependencies (it only handles these exact three *wh*-dependencies). So, this uncompressed explanation doesn't seem useful. In contrast, the explanations in (2b)-(2d) seem less cumbersome because they don't have to list out the individual dependencies, and they can generalize to other *wh*-dependencies.

(2) Possible explanations/theories
- a. Uncompressed:
  - •*Who* is allowed to appear at the front of a dependency that continues as *did she think the kitty was for* $\_\_{who}$
  - •*What* is allowed to appear at the front of a dependency that continues as *did she think* $\_\_{what}$ *was pretty*
  - •*Who* is not allowed to appear at the front of a dependency that continues as *did she think the kitty for* $\_\_{who}$ *was pretty*
- b. Compression 1: *Wh*-dependencies can't cross specific structures (known as syntactic islands: Ross 1967), such as a complex subject like $_{\textit{Subject-NP}}$*[the kitty* $_{PP}$*[for]]*, and many other structures.
- c. Compression 2: *Wh*-dependencies can't cross more than two specific structures (known as bounding nodes: Chomsky 1973, Huang 1982, Lasnik & Saito 1984), which are

fewer in number than syntactic islands.

d. Compression 3: *Wh*-dependencies can't cross too many low-probability structures (syntactic chunks of different sizes: Pearl & Sprouse 2013b, Dickson et al. 2022), which are more general-purpose than bounding nodes.

So, hopefully we can see that theories help us usefully compress the data we want to explain, and one key way theories do so is by utilizing key factors that underlie the observable data. If we think about theories this way, one key thing we're interested in when building a theory of the acquisition process is the set of relevant factors for explaining the acquisition process. So then, for building theories of language acquisition, what *are* all those relevant factors that we should keep in mind?

I'll begin by trying to answer this question about relevant factors, drawing on quantitative approaches to language acquisition– in particular, mathematical learnability and computational modeling – that concretely articulate theories of acquisition (or at least relevant parts of theories). I'll then try to synthesize some of the results coming from these approaches, prefaced by cautions and caveats for how to interpret those results, given the factors that different approaches tend to incorporate into their investigations. For our purposes, the insights gained through careful interpretation of quantitative investigations are (hopefully) some pieces of a complete, elegant theory of syntactic language acquisition.

2 What are the relevant factors?

The factors I want to focus on here come from theorists taking inspiration from behavioral data about both language development and general cognitive development in children. Figure 1 was a recent attempt I made (Pearl, 2023a) to highlight some factors that might be relevant from this perspective, organized into different types of factors (e.g., external vs. internal factors).

Figure 1 is meant to capture a snapshot of how acquisition iteratively unfolds over time. So, in effect, the proposal in Figure 1 specifies what can go on during one "unit" of time of the child's **learning period**, which is the time when acquisition occurs. The learning period itself could occur over days, months, or years, depending on the specific linguistic knowledge we focus on. With that in mind, each moment of time involves both external and internal components.

An external factor: Input. External components are observable. For instance, we can observe the **input** signal children encounter (e.g., the child language interactions they experience). The input signal is the physical signal in the world, such as auditory components like pitch and loudness of the utterance. For example, consider this utterance: "Why did you think it was a good idea to draw on the wall with permanent marker when Mommy wasn't paying attention?" This utterance might be said with a rising pitch contour and increasing volume. The input can also include other aspects of the environment, such as who said the utterance (e.g., an exasperated Mommy), where they said it (e.g., in the kitchen), when they said it (e.g., late morning), and what people or objects were in the environment at the time (e.g., Mommy, children, permanent markers, a marked-up wall, etc.).

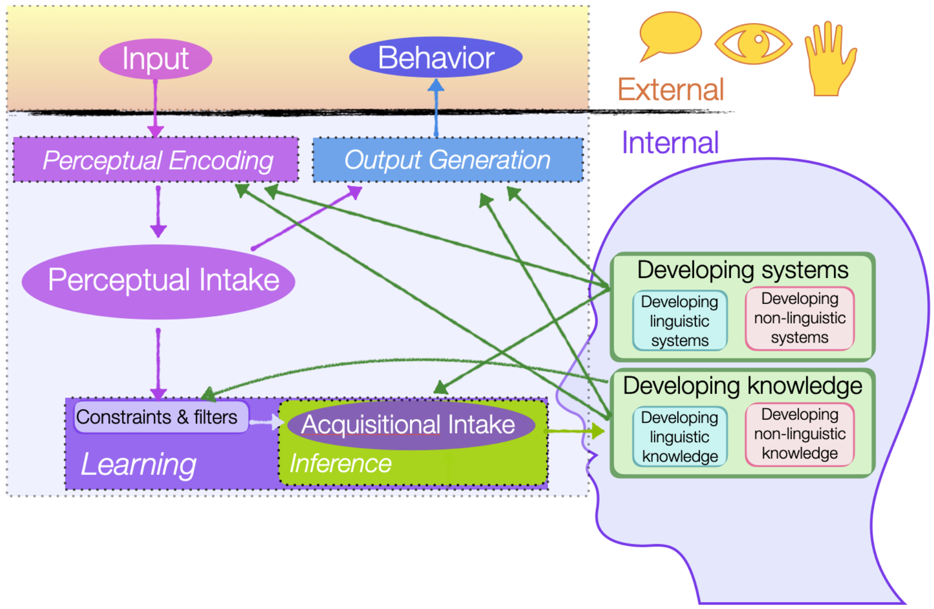

Figure 1: Some relevant parts of the acquisition process, as specified by Pearl (2023a). External components (input and behavior) are observable. Internal components aren't observable, and include perceptually encoding information from the input signal (yielding the perceptual intake), generating output from the encoded information (yielding the observable behavior), and learning from the encoded information (using constraints & filters to yield the acquisitional intake, and doing inference over that intake). The developing systems and developing knowledge (both linguistic and non-linguistic) impact all internal components, while the learning component updates the developing knowledge.

From external to internal: Perceptual encoding. One key transformation is from external to internal: from the observable input to the unobservable **perceptual intake**, which is the information the child is able to perceive and extract from the input signal. This perception and extraction process is the **perceptual encoding** that generates the perceptual intake from the input. Perceptual encoding draws on the child's developing knowledge and systems to extract information. One key aspect of perceptual encoding is that the child may imbue the input with information from her developing linguistic knowledge – that is, she "perceives" information that may not be explicitly in the input signal.

For instance, in our example utterance, the child may be able to perceive syllables (e.g., /waj/, /dɪd/, /ju/, /θɪŋk/ etc.), words (e.g., *why*, *did*, *you*, *think*, etc.), and syntactic structure (e.g., [$_{CP}$ Why did [$_{IP}$ you [$_{VP}$ think __$_{why}$ ...]]]). All this information comes from the child's developing linguistic knowledge, rather than being observable directly in the input. Other information may also be taken in, such as properties of the events described (thinking, drawing on the wall, paying attention), among many other types of information.

Under this proposal, what children can perceive depends on several things. First, children's perceptual intake depends on what they currently know about their language, given their developing linguistic knowledge (e.g., *why*, *think*, and *draw* are words; *why* is understood in a position later than than the position it's uttered in, etc.). Second, children's perceptual intake depends on what they currently know about the world, given their developing non-linguistic knowledge (e.g., what event Mommy is probably asking why about, given that she's not at all pleased). Third, children's perceptual intake depends on how well they can extract information of different kinds, given their developing linguistic systems (e.g., speech segmentation, syntactic parsing, *wh*-dependency resolution) and their developing non-linguistic systems (e.g., memory, cognitive inhibition). Notably, extracting information from the input signal involves ignoring information present (e.g., where the utterance was spoken) and adding information not explicitly present (e.g., what the words are in the speech stream, how words group together to form constituents). What children ignore and add to generate their perceptual intake depends on their developing knowledge and developing systems.

What to learn from: The acquisitional intake. We then get to the **acquisitional intake**, which is the information that children actually learn from. At this point, there's an additional key transformation. The basic idea is that children don't learn from all the information they can perceive. So, for instance, even if they can perceive all the words in the utterance from before, it may be less relevant that individual words are present in the speech stream (e.g., *why*, *did*, *you*, *think*) if children are fairly confident about their speech segmentation. Instead, if children are still learning about the distribution of *wh*-dependency structures in their language, what may be more relevant are characteristics of the *wh*-dependency in that utterance. These characteristics include the structures the words are part of (e.g., CP, IP, VP, etc.) and where the *wh*-word *why* is interpreted (e.g., in the VP with *think* as opposed to the VP with *draw* or the VP with *paying attention*). This is where the **constraints & filters** of Figure 1 come in: focusing the child on the relevant information in the perceptual intake, given the current stage of acquisition.

Importantly, what's perceived as "relevant" depends on what children are trying to learn, and what hypotheses they're considering. In terms of the proposal in Figure 1, constraints can help define what hypotheses are worth considering (that is, how the child's hypothesis space should be usefully constrained: Heinz & Rawski 2024). For instance, in our *wh*-question example, suppose the child is learning about the distribution of *wh*-dependencies. Perhaps her hypothesis space is then constrained to the set of possible *wh*-dependencies and their relative frequencies, characterized by the syntactic structures that comprise those dependencies. Then, given that hypothesis space, only certain aspects of the perceptual intake become relevant for the utterance above: specifically, where "why" is intended to be interpreted (in the VP with *think* as in [$_{CP}$ Why did [$_{IP}$ you [$_{VP}$ think $\_\_{why}$ ...]]]).

Filters may also focus the child's attention on a subset of the perceptual intake, irrespective of any hypothesis space constraints. For example, an attentional filter might cause the child to mis-interpret the *why* because her attention is drawn to emotionally-charged components of the utterance (e.g., *draw on the wall...*). That is, she might mis-interpret *why* as being understood in the VP with *draw*, rather than in the VP with *think*, even though she perceived the VP with *think* as part of her perceptual intake. An interesting effect of mis-interpreting *why* (whatever the cause) is that the acquisitional intake in this case is actually a skewed version of what an adult would extract from the original input (i.e., a *wh*-dependency with *why* interpreted somewhere other than with the

VP containing *think*). That is, what the child is learning from is different than what seems to be available in the input signal, at least according to adult perceptions.

This is not to say that the acquisitional intake is always skewed like this, compared to the "intended" intake. It could be that the child in the example above would interpret *why* as an adult would despite any constraints or filters, and so her acquisitional intake would be similar to an adult's. Still, the simple fact that the available signal the child learns from might indeed be transformed is worth noting in any theory of the acquisition process because the child's acquisitional intake could, indeed, be different from an adult's. Whatever the resulting information in the acquisitional intake, that's the information the child uses to update her internal state (in this example: to update her hypotheses about *wh*-dependency distributions).

The learning part: Inference. The process of updating on the basis of the acquisitional intake is captured in the "learning" block of Figure 1, under **inference**. In particular, this internal piece concerns how the child's developing knowledge (both linguistic and non-linguistic) is updated over time. Inference typically involves non-linguistic abilities that themselves may be developing, like probabilistic inference, statistical learning, or hypothesis testing. The result of this inference can be used to update the developing knowledge – potentially both linguistic knowledge and non-linguistic knowledge. For instance, in our *wh*-dependency example, the child might update her hypotheses about how likely *why* is to be interpreted with *think* as a main verb (linguistic knowledge) and how likely adults like Mommy are to ask about why the child thought something was a good idea (non-linguistic knowledge).

From internal to external: Output generation and observable behavior. Whatever the contents of the child's developing knowledge, that is what she uses to generate the observable **behavior** we see externally as output. Key to this relationship between the internal knowledge of the child and her observable behavior is the **output generation** of Figure 1. In particular, output generation depends both on the child's current perceptual intake, as well as her developing knowledge and developing systems.

For example, suppose we're trying to elicit a child's knowledge about *wh*-dependencies, and ask for her interpretation of "Where did Lizzie say she was going to catch butterflies?" (Omaki et al., 2014). Her perceptual intake is whatever information she can encode from that utterance, and then, to generate observable behavior, she draws on her developing knowledge (e.g., linguistic: about *wh*-dependency distributions in her language and conversational goals; non-linguistic: about where Lizzie is likely to do different actions) and her developing abilities (e.g., linguistic: utterance generation; non-linguistic: motor control, attention, and decision-making).

As an example of how developing knowledge can impact output generation, let's consider children's pragmatic knowledge, which allows them to understand the conversational intent behind any particular utterance. This knowledge is typically still developing by five years old (Papafragou & Musolino, 2003, Musolino & Lidz, 2006). As a specific example of pragmatic knowledge, children are sensitive to the implied question that an utterance corresponds to (sometimes called the Question Under Discussion: **QUD**) (Gualmini, 2004, Gualmini et al., 2008, Di Bacco et al., 2017, Savinelli et al., 2017, 2018, Scontras & Pearl, 2021). If the pragmatic context for an utterance isn't sufficiently supportive of the appropriate QUD, children can struggle to display other underlying linguistic knowledge.

In our *wh*-dependency example ("Where did Lizzie say she was going to catch butterflies?"), a child (like an adult) might normally interpret the question as asking about where the butterfly-catching happened (e.g., in a field). This is especially true if the context of the question was something that focused the conversational topic on butterfly-catching (e.g., "Lizzie likes catching butterflies in different places."). A potential QUD that might result from this context is "*Where does Lizzie like catching butterflies?*" Asking about where Lizzie was going to catch butterflies follows naturally from this QUD.

However, what if we wanted to know whether children can get the interpretation that asks about the location of the saying event (e.g., at home)? The previous context and resulting QUD focused on butterfly-catching seem to make it hard to lead to this interpretation. Adults may still demonstrate that they can get the saying-location interpretation, but children will likely struggle to. However, suppose the pragmatic context suggests that the QUD is about the saying event (e.g., "I'm trying to remember where Lizzie was saying something about butterflies." $\rightarrow$ Potential QUD: *Where did Lizzie say something about butterflies?*). Given this QUD, children might be better able to interpret the question ("Where did Lizzie say she was going to catch butterflies?") as asking about where Lizzie did the saying. So, we can hopefully see from this example how children's developing pragmatic knowledge about QUDs may impact their output generation about interpretations for *wh*-dependencies.

Let's now come back to the generated output that signals the underlying knowledge – that is, the observable behavior. We can observe things like linguistic productions or behavior (either naturalistic or coming from clever experimental designs that elicit those productions or behavior). The behavior could be linguistic (e.g., answering where Lizzie did the saying vs. answering where Lizzie did the butterfly-catching) or non-linguistic (pointing at a picture of the location where Lizzie was speaking vs. a picture of the location where Lizzie was catching butterflies). By seeing how the child interprets these *wh*-dependencies (e.g., where the saying happened vs. where the butterfly-catching happened), we can observe which *wh*-dependency she thinks is more likely in context; her preferred *wh*-dependency in this context demonstrates her internalized knowledge about the distribution of *wh*-dependencies in her language.

Back to acquisition theorizing. My goal in walking through the details of this proposal about the acquisition process was to highlight both (i) the factors that we should consider when building our acquisition theories, and (ii) how complicated it can get. One striking thing (to me) about Figure 1 is how connected everything is (i.e., there are a lot of arrows). We might reasonably wonder if we can make any progress at all on theorizing with this kind of interconnectivity present – that is, until we know for sure about factor X (e.g., developing systems), how can we possibly say anything sensible about factors Y and Z (e.g., perceptual encoding, inference) that depend on factor X?

In response to this question, I (currently) feel that quantitative approaches like mathematical learnability and computational modeling have the right idea: we have to idealize somewhere to get anywhere. Perhaps we ignore (for now) the impact of certain factors. Or, perhaps we simplify the contributions of any particular factor so that we can actually implement (part of) a theory concretely and evaluate that theory (part). Or, perhaps we ignore the potential impact of "noise" on any particular factor, such as noise in the child's input from accidental mis-speakings, noise in the child's perceptual encoding due to misparsing, or noise in the child's inference due to a skewed update calculation. We'll still learn something about how good the implemented theory (part) is — or isn't – for explaining the acquisition data we're trying to explain.

With that said, I do think it's useful to know what might be considered relevant factors even while we might put some aside for the moment, because there will (hopefully) come a time when we won't need to idealize so much. In the meantime, we can at least implement what we can of the factors we think are relevant, thereby implementing theories of the acquisition process that we can then evaluate against the available empirical data.

3 What we've learned so far from quantitative approaches

Now I'll review some of the insights we've gained from different quantitative approaches, like mathematical learnability and computational modeling. I'll begin first with some caveats about interpreting results from these approaches, discuss some of the factors different quantitative approaches tend to consider relevant, and then survey results from quantitative approaches, keeping in mind how they help us build acquisition theories.

3.1 Caveats: What can they tell us anyway?

When interpreting the results of quantitative investigations, it's important to remember what it is they're actually investigating – in particular, which aspect(s) of the acquisition process they're implementing concretely and evaluating. Moreover, any acquisition theory implemented by a quantitative approach is a "(dis)proof of concept": the modeled child operating in the acquisition scenario formalized by that quantitative approach will behave in the predicted way (Pearl, 2021, 2023a,b, Portelance & Jasbi, 2024). Hopefully, that predicted way connects in some interpretable way to children's behavior, so that we can evaluate whether *that implementation of the acquisition theory* works (or doesn't). This is an important point: we can only interpret the results with respect to that acquisition theory, as implemented by the quantitative approach. If there are other theories (or other versions of the implemented theory), the results *don't* apply – they only apply to the implemented acquisition theory. This is why the assumptions about how relevant factors are implemented (or ignored) are so important: these are the assumptions that implement the acquisition theory being evaluated. It is only this acquisition theory that any results can then provide evidence for (or against).

With these considerations in mind, an acquisition theory can try to provide an explanation at different levels in the sense of Marr (1982): computational, algorithmic, and implementational (Pearl, 2023b). In my experience, computational-level acquisition theories are committed to the specified mental computations being performed (e.g., Bayesian inference over a certain acquisitional intake), but not necessarily committed to actual children performing them the way the modeled child does. Ideal/rational learner models (e.g., Foraker et al., 2009, Hsu & Chater, 2010, Pearl, 2011, Perfors et al., 2011, Hsu et al., 2011, Feldman et al., 2013, Hsu et al., 2013, Orita et al., 2013, Abend et al., 2017, Pearl et al., 2017, Nguyen & Pearl, 2019, Pearl & Sprouse, 2019, 2021, Dickson et al., 2022, Pearl & Forsythe, 2022, Dickson et al., 2024) as well as investigations leveraging large language models (e.g., Wilcox et al., 2023, Yedetore et al., 2023, Lan et al., 2024) typically take this approach, and so explain what is (im)possible to do, given the acquisition scenario specified.

Algorithmic-level acquisition theories are committed to the specified mental computations being performed using the specific steps the modeled child uses (e.g., a step-by-step approximation of Bayesian inference over a certain acquisitional intake received incrementally over time). Constrained/process learner models (e.g., Regier & Gahl, 2004, Yang, 2004, Pearl & Lidz, 2009, Pearl & Sprouse, 2013a, Pearl & Mis, 2016) typically take this approach, and so explain what is (im)possible for children to do, given the acquisition scenario specified.

Implementational-level models are committed to the specified mental computations being performed using the specified steps within a specified medium (e.g., a collection of neurons organized a particular way in a brain). To my knowledge, we don't currently have acquisition theories of this kind, but if we did, they would explain what is (im)possible for children's brains to do, given the acquisition scenario specified.

Ideally, if we had an explanation for how acquisition worked at two different levels (e.g., computational and algorithmic), we would be able to directly link them and so have a more complete explanation. For instance, a computational-level explanation would explain the acquisition computation that actually occurs, and the algorithmic-level explanation would explain the steps that get carried out. We would then be able to link these together in order to understand how the steps of the algorithmic-level explanation carry out the acquisition computation of the computational-level explanation. That is, we would have a theory of what's being computed during acquisition, and what cognitively-plausible steps are used to do that computation. In practice, it's not always obvious how to link explanations together like this. In the meantime, we can at least work on coming up with explanations at whatever levels we can, with the eventual goal of linking them up.

Explanations targeting different levels naturally tend to focus on some of the available relevant factors, rather than others. For instance, a computational-level theory may not be concerned with the learning period and developing non-linguistic abilities that impact inference, while an algorithmic-level theory may well be. I'll now attempt to briefly highlight factors that different quantitative approaches have focused on, as well as what seem to me to be some of the questions these approaches have focused on. By understanding the factors and questions of interest, we can better understand both (i) what level of explanation is being targeted, and (ii) what acquisition theory (part) is being implemented and evaluated by these quantitative approaches.

3.2 Relevant factors for different quantitative approaches

3.2.1 The factors that mathematical learnability approaches have focused on

Mathematical learnability approaches (e.g. Gold, 1967, Horning, 1969, Wharton, 1974, Angluin, 1980, Wexler & Culicover, 1980, Osherson et al., 1983, 1986, Valiant, 1984, Angluin, 1988b, Gibson & Wexler, 1994, Niyogi & Berwick, 1996, Osherson et al., 1997, Niyogi, 2006, Clark &

Eyraud, 2007, Hsu & Chater, 2010, Hsu et al., 2011, 2013, Clark & Lappin, 2011, Chater et al., 2015, Heinz, 2016, Rawski, 2021, De Santo & Rawski, 2022, Heinz & Rawski, 2024) have tended to focus on computational-level theories – that is, what is (im)possible for any learner to do, given a specific acquisition scenario. The acquisition scenarios they explore have typically highlighted the role of certain acquisition theory factors:

- the child's **acquisitional intake**
- the **constraints** that helpfully restrict the child's hypothesis space
- the **inference** process that leads to updating a child's developing linguistic knowledge
- the target **linguistic knowledge** state

So, mathematical learnability approaches have typically asked questions about how different factors impact the acquisition scenario, and what the subsequent impact is on learnability for that acquisition scenario. That is, the hallmark of mathematical learnability has been manipulating these factors in the ways described below (among others) and then seeing the result on what can (in principle) be learned.

For instance, when considering the child's **acquisitional intake**, is the child only learning from examples of what's in the language (direct positive evidence: Gold 1967, Angluin 1980, Valiant 1984, Angluin 1988b, Gibson & Wexler 1994, Niyogi & Berwick 1996, Niyogi 2006, Clark & Eyraud 2007, Hsu & Chater 2010, Hsu et al. 2011, 2013, Clark & Lappin 2011, Chater et al. 2015, Heinz 2016)? Or instead, is the child also getting feedback directly or indirectly about what isn't in the language (direct negative evidence: Gold 1967, Angluin 1980, Valiant 1984, Angluin 1988b, Heinz 2016; indirect negative evidence: Clark & Lappin 2011)? Is the child able to actively learn by asking specific questions about what she wants to know (targeted queries: Angluin 1988b, Heinz 2016)? What happens if the intake is "messy", containing items (perhaps by accident) that aren't actually in the target language (noise in the intake: Horning 1969, Clark & Lappin 2011, Heinz 2016)? What kind of information is available anyway (e.g., semantic information along with syntactic structure when learning about syntax: Wexler & Culicover 1980)?

When thinking about **constraints** on the child's hypothesis space, what kind of prior knowledge is necessary to helpfully restrict that hypothesis space (Heinz & Rawski, 2024)? When investigating how the **inference** process might work, do the "learning algorithms" that lead to acquisition success involve some element of randomness (e.g., guessing: Gold 1967; sampling: Niyogi & Berwick 1996)? Do successful algorithms involve sensitivity to errors when predicting the input (error-driven learning: Wexler & Culicover 1980, Gibson & Wexler 1994, Niyogi & Berwick 1996, Yang 2002, Heinz & Rawski 2024)?

For the child's acquisition target, does the internalized **linguistic knowledge** concern the entire language system (e.g., grammar: Gold 1967, Horning 1969, Wexler & Culicover 1980, Gibson & Wexler 1994, Niyogi & Berwick 1996, Niyogi 2006, Heinz 2016, De Santo & Rawski 2022, Heinz & Rawski 2024) or instead certain key components of the language system (e.g., a transformational component: Wexler & Culicover 1980)? Does the internalized knowledge have to be exactly right, or instead "close enough" to satisfy some other objective (Wharton, 1974, Valiant, 1984, Clark & Lappin, 2011)?

In contrast to computational-level theories exploring just these factors, computational modeling has provided a complementary set of explorations, investigating both computational-level and algorithmic-level theories. Computational modeling has also included other factors that mathematical learnability approaches have typically idealized away from, such as realistic input, a cognitively-plausible learning period, observable behavior, and the impact of both developing knowledge and developing systems.

3.2.2 The factors that computational modeling approaches have focused on

A useful distinction. I first want to distinguish between two types of computational modeling: cognitive and non-cognitive. This isn't necessarily a distinction made overtly in the computational modeling field yet, but I feel it's a helpful distinction. By "cognitive", I mean a model (i.e., a modeled learner) that takes cognitive plausibility seriously in its implementation. For instance, a cognitive modeled learner may extract the acquisitional intake from child-plausible input, have a child-plausible learning period length, and try to match observed child behavior. That is, the acquisition scenario that a cognitive modeled learner is implementing and evaluating is easily interpretable with respect to actual children. Moreoever – and this is most important for our purposes here – the acquisition theory implemented by a computational cognitive model tends to be easily interpretable: we can transparently see how the different acquisition theory factors are implemented and how those implemented factors impact the resulting modeled acquisition process. This has been the approach of most computational models in the acquisition literature (e.g., most of the ones I mentioned in section 3.1) until the advent of large language models (**LLMs**).

To me, LLMs are computational non-cognitive modeling, at least as implemented currently. A non-cognitive modeled learner maybe learns in a human-like way...but maybe it doesn't. To my knowledge, the learning mechanism underlying current successful LLMs is still being (very actively!) investigated. That is, we know the mechanics of how LLMs learn (i.e., a certain way of updating vectors of numbers on the basis of input, with a certain goal in mind, such as predicting the next word in a sequence). Yet, we don't necessarily understand how those step-by-step mechanics lead to the interpretable behavior we see the LLMs producing. In other words, LLMs use an algorithm, but it's currently unclear if that algorithm is intended to be an algorithm that humans would use. So, because the steps that LLMs are following are not intended to be the steps human children would follow, LLMs are not implementing an algorithmic-level acquisition theory.

Moreoever, it's unclear that many LLMs are implementing a computational-level acquisition theory either: LLMs at the moment don't tend to learn from child-plausible input or have a child-plausible learning period length, though they may try to match observed child target behavior. That is, the acquisition scenario that LLMs investigate can be difficult to link to the acquisition scenario that children face. Most importantly, many acquisition theory factors are hard to "read off" from current LLMs. What is the acquisitional intake? How did the developing knowledge impact the process of transforming the input into the acquisitional intake? Were there constraints on the learner's hypothesis space – and if so, what were they? These are questions that a computational-level acquisition theory would aim to answer. To my knowledge, LLMs don't answer these questions in a transparent way, and so I don't believe we should call them cognitive models.

So, given that LLMs are often not computational cognitive models, why should we bother with them at all if we're interested in acquisition theory-building? My current feelings align with those of Futrell & Mahowald (2025), who discuss the Contravariance Principle (Cao & Yamins, 2021): "[I]f a computational problem is hard, then there are only a few ways to solve it, so we expect different systems (for example, neural networks and the brain) to converge to the same solution".

As an example, vision is a hard computational problem, and it turns out that there are striking parallels in the way that humans, monkeys, non-primate animals, and neural networks "solve" vision (Rajalingham et al., 2018, Pungor et al., 2023). Similarly, language acquisition is a hard computational problem. Human children have solved this problem. LLMs are modeled learners implemented with neural networks who seem able to learn something that allows them to generate sophisticated linguistic behavior. The important question is this: does that sophisticated linguistic behavior mean that LLMs have "solved" language acquisition too?

I'm not currently convinced it does (see Lan et al. 2024 for similar reservations). One concern for me is identifying when performance counts as "close enough" to human language performance. For instance, is judging one item as better (more grammatical) than another 80% of the time good enough when humans make that judgment 99% of the time? How do we decide? If an LLM reaches the "good enough" threshold for one item, is that enough? In other words, how many items have to have good enough performance? Is having good enough performance on one set of related items (e.g., *wh*-dependency knowledge) "good enough" if the same LLM doesn't have good enough performance on some other set of items (e.g., agreement morphology, binding relations, passives, etc.)? Does it matter if LLMs can learn languages that seem "impossible" for humans to learn (Yang et al., 2025)? That is, if the LLMs can learn these impossible languages but humans can't, are the LLMs not "close enough" to human language performance (in this case, by being "better" than humans at learning seemingly-impossible languages)? Answering these questions will help us identify when we have an LLM that has "solved" language acquisition.

To the extent that we believe any given LLM's sophisticated linguistic behavior indicates that the LLM has "solved" language acquisition, then we might expect both humans and those LLMs to "converge to the same solution" for solving language acquisition. Any such convergence between humans and LLMs would be inspiration for how key components of the human acquisition process work.

So, with all that said, I believe non-cognitive computational models like LLMs have the potential to inform acquisition theory if implemented in particular ways that align with the acquisition scenarios of human children. Below, I'll discuss some ideas about how such LLMs have been implemented. Still, the key point is that thoughtfully-implemented LLMs, though themselves non-cognitive, might be leveraged to answer cognitive questions, which are the type of questions that an acquisition theory cares about.

Factors that computational cognitive modeling has focused on. Computational cognitive models, as mentioned above, seem to more naturally lend themselves to acquisition theory-building. The ones I'm aware of typically implement a theory at the computational level (e.g., ideal/rational modeled learners: Foraker et al. 2009, Hsu & Chater 2010, Pearl 2011, Perfors et al. 2011, Hsu et al. 2011, Feldman et al. 2013, Hsu et al. 2013, Orita et al. 2013, Abend et al. 2017, Pearl et al. 2017, Nguyen & Pearl 2019, Pearl & Sprouse 2019, 2021, Dickson et al. 2022, Pearl & Forsythe

2022, Dickson et al. 2024) or the algorithmic level (e.g., constrained/process modeled learners: Regier & Gahl 2004, Yang 2004, Pearl & Lidz 2009, Pearl & Sprouse 2013a, Pearl & Mis 2016). They often implement components corresponding to these acquisition theory factors:

- the child's **learning period** (often converted into the quantity of input a child would encounter during that time)
- the child's **acquisitional intake** and the related **constraints** and **filters** defining it,
- the **inference** process that leads to updating a child's developing linguistic knowledge
- the target **linguistic knowledge** state and target **observable behavior**
- the impact of the child's **developing knowledge** on the child's intake and observable behavior
- the impact of the child's **developing systems** on the child's intake and observable behavior

As with mathematical learnability approaches, computational cognitive models have typically asked questions about how different factors affect acquisition success. More specifically, does the acquisition theory implemented the way it is in the modeled learner allow the learner to succeed (e.g., generate the target observed behavior, given plausible child language input and a plausible learning period)? For instance, here are some questions different computational cognitive models have explored:

- How specific do the **constraints** on children's hypothesis space need to be that define the **acquisitional intake** (e.g., Pearl & Sprouse, 2013a,b, Pearl, 2014, Pearl & Sprouse, 2015, Pearl, 2017, Bates & Pearl, 2019, Pearl & Bates, 2022, Dickson et al., 2022, 2024)?
- What information should be included vs. **filtered** out in children's acquisitional intake (Reali & Christiansen, 2005, Pearl, 2007, Pearl & Weinberg, 2007, Pearl & Lidz, 2009, Pearl & Mis, 2011, Perfors et al., 2011, Frank et al., 2013, Orita et al., 2013, Pearl & Sprouse, 2013a,b, Pearl, 2014, Pearl & Sprouse, 2015, Pearl & Mis, 2016, Abend et al., 2017, Fitz & Chang, 2017, Nguyen & Pearl, 2018, Bates & Pearl, 2019, Nguyen & Pearl, 2019, 2021, Dickson et al., 2022, Pearl & Bates, 2022, Maitra & Perkins, 2023, Perkins & Hunter, 2023, Dickson et al., 2024)? How might such a **filtering** process arise in the first place (Perkins et al., 2022, Maitra & Perkins, 2023, Perkins & Hunter, 2023, Perkins et al., 2024)?
- How might children's uncertainty about the true (target) linguistic representation for an input item (due to their **developing knowledge** and **developing systems**) affect when they include an item in their acquisitional intake (Fodor, 1998a,b, Sakas & Fodor, 2001, Sakas & Nishimoto, 2002, Yang, 2002, 2004, Fodor & Sakas, 2005, Fodor et al., 2007, Legate & Yang, 2007, Pearl, 2007, Pearl & Weinberg, 2007, Foraker et al., 2009, Pearl, 2009, Pearl & Mis, 2011, Sakas & Fodor, 2012, Yang, 2012, Pearl & Mis, 2016, Sakas, 2016, Fodor, 2017, Fodor & Sakas, 2017)?

- How are existing successful theories affected by naturally-occurring variation in children's input and acquisitional intake – do these theories still succeed (Bates & Pearl, 2019, Pearl & Bates, 2022)?

- When is children's **observable behavior** due to still-developing knowledge vs. still-developing systems (Freudenthal et al., 2007, 2009, 2010, 2015, Savinelli et al., 2017, 2018, Forsythe & Pearl, 2020, Scontras & Pearl, 2021, Pearl & Forsythe, 2022, Freudenthal et al., 2024)?

Computational cognitive models contrast with non-cognitive models because, as mentioned above, the non-cognitive models don't transparently implement most acquisition theory factors. With that said, I review ways that I believe non-cognitive models can still be useful.

Factors that computational non-cognitive modeling has focused on. Computational models that are non-cognitive, like LLMs, have recently been used to investigate the relationship between the information in the input children encounter and the ability to develop human-like syntactic knowledge from that input (e.g., Wilcox et al., 2023, Yedetore et al., 2023). More specifically, LLMs are a way to explore the information that's in principle available to children. LLMs do this by being very good learners. So, the idea is that if an LLM can learn the appropriate knowledge given the available input, that knowledge is in principle learnable, given whatever built-in knowledge the LLM has. Then, we care about what that built-in knowledge is. In other words, using LLMs to assess the information available in the input can inform a key factor of an acquisition theory: the prior knowledge that's necessary so that children can in fact develop the appropriate linguistic knowledge from their available input (see Pearl (2022) for more discussion of this point). The LLMs I'm currently aware of that may inform acquisition theory typically focus on these acquisition theory factors:

- the **input** to acquisition, which ideally would correspond to a child's input (e.g., Yedetore et al., 2023). However, current LLMs often require training on very large corpora (e.g., GPT-3 required data equivalent to over 10,000 years of linguistic experience: Lan et al. 2024), and these corpora often aren't child language interactions.

- the **constraints** and **filters** that yield successful learning (though these factors are notoriously hard to interpret directly from an LLM)

- the target **observable behavior**, which ideally would correspond to a child's target **linguistic knowledge**.

As I mentioned above, non-cognitive computational models typically focus on questions about the information available in principle in the child's **input**, and what must be built in (underlying **constraints** or **filters**) to achieve the target **observable behavior** from that input.

3.3 Insights about relevant factors

When we consider how to interpret the results of quantitative investigations, it's useful to be specific about which acquisition theory factors the result speaks to. With this in mind, I'll discuss factors we've gained insight about, and what those specific insights are, given the results from different quantitative approaches.

3.3.1 Intake, constraints, and developed linguistic knowledge

Interestingly (and perhaps unsurprisingly given the interconnectivity demonstrated in Figure 1), several insights involve the intake, constraints, and developed linguistic knowledge considered together. That is, if the intake is like $Intake_A$, then the necessary constraints are $Constraints_A$ and the developed knowledge is $DevelopedKnowledge_A$. If instead, the intake is like $Intake_B$, then the necessary constraints should be $Constraints_B$ and the developed knowledge should be $DevelopedKnowledge_B$. From mathematical learnability investigations (Gold, 1967, Horning, 1969, Wiehagen, 1977, Angluin, 1988b, Heinz, 2016, De Santo & Rawski, 2022), we find that children's hypothesis space (and therefore the possible target knowledge children aim to develop) must be helpfully constrained for any chance of success, given plausible assumptions about children's intake and the nature of the acquisition task (e.g., children's limited learning period and their developing cognitive systems). In particular, a collection of investigations (Gold, 1967, Berwick, 1985, Angluin, 1988b, Heinz, 2007, Tîrnăucuă, 2008, Heinz, 2009, Clark & Lappin, 2011, Heinz, 2016) suggests that if children's intake is restricted to "positive evidence" (only examples of what's in the language, rather than also pointing out what's not), then human languages must be more constrained than we might otherwise expect in order for acquisition to succeed.

Recent investigations with LLMs suggest some particular ways that human languages may be constrained, because sufficiently powerful learners (like LLMs and potentially humans) find it easier to learn languages with these properties. One property involves the idea of "information locality" (Kallini et al., 2024), where "elements that statistically predict each other are close to each other" (Futrell & Mahowald, 2025). In other words, the hypothesis space children have for human languages may be biased towards languages that have this information locality property.

Another property involves the observation that the information conveyed by word order in human languages is often signaled by other linguistic cues like meaning or case markers (Futrell et al., 2015, Koplenig et al., 2017, Pijpops & Zehentner, 2022, Mahowald et al., 2023, Futrell & Mahowald, 2025). That is, word order is somewhat redundant in many cases and so powerful learners (like some LLMs) can be "relatively insensitive" (Futrell & Mahowald, 2025) to word order, and still learn appropriate linguistic knowledge. This finding suggests that the hypothesis space children have for human languages may be biased towards languages that allow some word order flexibility.

More generally, recent LLM-based investigations suggest that children's hypothesis spaces, while needing constraints or biases, may not need the specific constraints traditionally proposed in the learnability literature. That is, the implementation of the necessary constraints may be different than we originally thought, though children do in fact need constraints on their hypothesis space. Moreover, LLMs may offer a new answer for how certain constraints on the hypothesis space arise. The idea is that a "prohibited" representation is still in the learner's hypothesis space. That is, it is possible in principle for the learner to learn that prohibited representation–but something prevents that prohibited representation from actually being learned. Prior computational cognitive modeling work (Perfors et al., 2011) suggested that the "something" was a built-in general-purpose bias for the simplest representation that still accounted for the input. Recent LLM investigations suggest a different possibility: the way learning itself works in some sufficiently powerful learners (e.g., Transformer-based LLMs: Hahn & Rofin 2024) prevents the learner from actually converging on the prohibited representation. So, under this LLM-based proposal, hypothesis space constraints emerge from how learning works, not because there's an inherent representational bias defining the hypothesis space.

Investigations using computational cognitive modeling align with this idea that hypothesis space constraints, however they arise, may be different from what acquisition theorists traditionally thought they were. More specifically, some constraints may be more "general-purpose", targeting not just (the exclusion of) specific linguistic elements but instead targeting the shape of the allowed linguistic system more broadly. For instance, to learn certain restrictions on *wh*-dependencies in English (known as "syntactic islands"), children's hypothesis space doesn't need to be defined in terms of certain linguistic structures ("bounding nodes") that only arise when considering if a *wh*-dependency crosses a syntactic island (Chomsky, 1973, Rizzi, 1982, Huang, 1982, Lasnik & Saito, 1984, Torrego, 1984). Instead, children's hypothesis space can be defined more generally in terms of the structures that make up any *wh*-dependency (and indeed, any utterance) in the language (e.g., verb phrases and complementizer phrases) (Pearl & Sprouse, 2013a,b, Pearl, 2014, Pearl & Sprouse, 2015, Pearl, 2017, Bates & Pearl, 2019, Pearl & Bates, 2022, Dickson et al., 2022, 2024).

Still, mathematical learnability investigations also suggest that fewer constraints on children's hypothesis space may be necessary if in fact children's intake is not simply the direct positive evidence. For instance, acquisition success in a less-constrained hypothesis space is possible if children also get negative evidence (i.e., signals about what's not in the language) (Gold, 1967, Angluin, 1988b, Tîrnăucuă, 2008, Clark & Lappin, 2011). On the flip side, computational cognitive modeling investigations have suggested that helpful restrictions on the intake (i.e., learning only from a subset of the direct positive evidence) can lead to acquisition success, even within a less-constrained hypothesis space, just as some mathematical learnability investigations speculated (but didn't test at the time) (Gold, 1967, Heinz, 2016).

More specifically, if children filter out information present in the input – that is, their acquisitional intake ignores some of the positive evidence available – they can succeed at acquisition for several types of linguistic knowledge, including basic word order (Pearl, 2007, Pearl & Weinberg, 2007, Maitra & Perkins, 2023, Perkins & Hunter, 2023), English anaphoric *one* (Pearl, 2007, Pearl & Lidz, 2009), restrictions on *wh*-dependencies (Pearl & Sprouse, 2013a,b, 2015), passives (Nguyen & Pearl, 2019, 2021), and pronoun interpretation (Frank et al., 2013). This kind of filtering is particularly useful in realistic learning situations where noise is present, and children might be led astray if they naively tried to learn from all the available direct positive evidence. Interestingly, recent computational cognitive modeling work has suggested how children might learn what to filter out of their acquisitional intake, rather than needing to know specifically what to filter out beforehand (Perkins et al., 2022, Maitra & Perkins, 2023, Perkins & Hunter, 2023, Perkins et al., 2024). A key assumption is that children have "general-purpose" knowledge that some of the input signal is noise – not how much noise there is or what items are noise, but simply that there *is* noise that needs to be filtered out.

On the other hand, some investigations also find that fewer constraints on children's hypothesis space are necessary if children's intake expands to include a broader set of information, rather than contracting to a subset of direct positive evidence (Pearl, 2023a). Here are some highlights from recent computational cognitive modeling investigations:

- Including data from all utterances in the language when learning about English yes/no questions like "Was the penguin who was on the iceberg eating a fish?" (Reali & Christiansen, 2005, Perfors et al., 2011, Abend et al., 2017, Fitz & Chang, 2017)
- Including data from all *wh*-dependencies when learning about the acceptability of specific *wh*-dependencies (Pearl & Sprouse, 2013a,b, Pearl, 2014, Pearl & Sprouse, 2015, Pearl, 2017, Bates & Pearl, 2019, Pearl & Bates, 2022, Dickson et al., 2022, 2024)
- Including data from other pronouns besides *one* when learning about English anaphoric *one* (Pearl & Mis, 2011, 2016)
- Including lexical semantic information when learning about the syntax of the English passive (Nguyen & Pearl, 2018, 2019, 2021)
- Including discourse information when learning about the syntax of English pronouns (Orita et al., 2013)

3.3.2 Inference

We also have some insights about how children update their internal developing knowledge on the basis of their acquisitional intake – that is, in some sense, how the heart of "learning" operates, using whatever inference mechanism children use. Notably, two insights about children's inference are tied to other aspects of the acquisition process. The first insight, coming from mathematical learnability, ties inference to hypothesis space constraints: being able to leverage probability information during inference doesn't help with acquisition unless there are also constraints on children's hypothesis space (Angluin, 1988a, Heinz, 2016). That is, being able to navigate a hypothesis space in a more sophisticated way (by leveraging available probabilistic information) doesn't help much in principle if the hypothesis space isn't helpfully constrained in the first place, however those constraints may arise.

A second insight comes from computational cognitive modeling work, and connects inference to children's acquisitional intake. In particular, learning from data perceived as ambiguous is very helpful, rather than (waiting for and) only learning from data perceived as unambiguous (e.g., word order: Fodor 1998a,b, Sakas & Fodor 2001, Sakas & Nishimoto 2002, Yang 2002, 2004, Fodor & Sakas 2005, Fodor et al. 2007, Pearl 2007, Pearl & Weinberg 2007, Sakas & Fodor 2012, Yang 2012, Sakas 2016, Fodor 2017, Fodor & Sakas 2017; English anaphoric *one*: Pearl 2007, Foraker et al. 2009, Pearl 2009, Pearl & Mis 2011, 2016; optional infinitives: Legate & Yang 2007). To be sure, updating the internal knowledge more confidently when the data are perceived as unambiguous is fine – but children likely also update even when the data are perceived as ambiguous, in order to learn what they do as quickly as they do. This seems especially true during the beginning of acquisition, when so much is uncertain and most data are likely perceived by children as ambiguous.

A third insight involves learning in the real world, where children's input naturally varies due to a variety of factors, including cultural background, birth order, family size, number of caretakers present, diversity of at-home vs. away-from-home experiences, and socioeconomic status, among many others. A recent finding from computational cognitive modeling suggests that at least some existing learning algorithms (which involve inference on the basis of the acquisitional intake) are robust to naturalistic input variation (i.e., learning about constraints on *wh*-dependencies, given different socioeconomic status backgrounds: Bates & Pearl 2019, Pearl & Bates 2022).

3.3.3 Developing systems and knowledge

A final set of insights relates to children's developing cognitive systems and developing knowledge. Interestingly, a recent insight from LLM-based approaches is that there can be sudden leaps in internalized linguistic knowledge (as measured by performance on language tasks meant to assess internalized knowledge) *without* any internal systems overtly changing (Futrell & Mahowald, 2025). That is, learners who have very powerful learning mechanisms (the way LLMs do and children might) may show sudden improvements after plateauing for awhile. This sudden improvement is sometimes called "double descent" (Nakkiran et al., 2021) or "grokking" (Power et al., 2022), and occurs even though nothing internal has changed for the learner (i.e., increased memory, more accurate inference, etc.). Why this occurs is currently unclear, but the simple fact that it *can* occur suggests that children's cognitive systems supporting language acquisition may not need to change (e.g., a qualitative change due to maturation) in order for qualitative updates to occur in children's acquired knowledge. Instead, some internal magic happens after some crucial amount of data has been learned from, and this magic is not necessarily due to a change in the learner's developing linguistic and non-linguistic systems. Perhaps some internal reorganization of the linguistic or non-linguistic knowledge occurs that enables a sudden leap in acquisition performance – but we simply don't yet know what that change is and why it occurs. In short, the knowledge must develop, but the systems don't need to in order for that knowledge to develop.

On the flip side, sometimes it is in fact the internal systems that need to develop, and not the internal knowledge. More specifically, children's linguistic knowledge may already be the target knowledge (i.e., adult-like), but the systems that support children demonstrating that knowledge may still need to develop. Insights from computational cognitive modeling suggest that, in some cases, non-adult-like observable behavior is due to still-developing cognitive systems rather than still-developing linguistic knowledge (e.g., optional infinitives: Freudenthal et al. 2007, 2009, 2010, 2015, 2024; quantifier scope resolution: Savinelli et al. 2017, 2018, Scontras & Pearl 2021; pronoun interpretation: Forsythe & Pearl 2020, Pearl & Forsythe 2022)

4 Conclusion

Our goal is to build an insightful, elegant theory of syntactic acquisition – that is, a useful abstraction that compactly explains the linguistic behavior we observe during acquisition. Such a theory will articulate the acquisition process by specifying the relevant factors and how they work together to generate the empirical data we have about language acquisition. Here I've tried to provide a sense of what we think those relevant factors might be and what we currently know about them, courtesy of quantitative approaches that use either mathematical learnability techniques or computational modeling techniques. These approaches allow us to precisely articulate a particular theory of the acquisition process (or at least key factors involved in an acquisition theory). Of course, there seem to be quite a few factors that are relevant, and they interact with each other. Nonetheless, these approaches have yielded insights about the nature of (at least some of) these factors and their interactions, when it comes to explaining children's syntactic acquisition. By continuing to use these methods, we can hopefully gain more insights that allow us to build more complete theories of how children learn what they do as fast and as well as they do.

References

Abend, O., Kwiatkowski, T., Smith, N. J., Goldwater, S., & Steedman, M. (2017). Bootstrapping language acquisition. *Cognition*, *164*, 116–143.

Angluin, D. (1980). Inductive inference of formal languages from positive data. *Information and Control*, *45*(2), 117–135.

Angluin, D. (1988a). *Identifying languages from stochastic examples*. Yale University. Department of Computer Science.

Angluin, D. (1988b). Queries and concept learning. *Machine Learning*, *2*(4), 319–342.

Bates, A. & Pearl, L. (2019). **What do you think that happens?* A quantitative and cognitive modeling analysis of linguistic evidence across socioeconomic status for learning syntactic islands. In Megan M. Brown and Brady Dailey (Ed.), *Proceedings of the 43rd annual Boston University Conference on Language Development*, (pp. 42–56)., Somerville, MA. Cascadilla Press.

Berwick, R. C. (1985). *The acquisition of syntactic knowledge*, volume 16. MIT press.

Cao, R. & Yamins, D. (2021). Explanatory models in neuroscience, Part 2: Functional intelligibility and the contravariance principle. *Cognitive Systems Research*, *85*, 101200.

Chater, N., Clark, A., Goldsmith, J., & Perfors, A. (2015). *Empiricism and language learnability*. Oxford University Press.

Chomsky, N. (1973). Conditions on transformations. In S. Anderson & P. Kiparsky (Eds.), *A Festschrift for Morris Halle* (pp. 237–286). New York: Holt, Rinehart, and Winston.

Clark, A. & Eyraud, R. (2007). Polynomial identification in the limit of substitutable context-free languages. *Journal of Machine Learning Research*, *8*(Aug), 1725–1745.

Clark, A. & Lappin, S. (2011). *Linguistic Nativism and the Poverty of the Stimulus*. Oxford: Wiley-Blackwell.

De Santo, A. & Rawski, J. (2022). Mathematical linguistics and cognitive complexity. In *Handbook of Cognitive Mathematics* (pp. 1–38). Springer.

Dennett, D. C. (1990). The interpretation of texts, people and other artifacts. *Philosophy and Phenomenological Research*, *50*, 177–194.

Di Bacco, F., Tieu, L., Sevdali, C., Folli, R., Moscati, V., & Romoli, J. (2017). Testing the QUD approach: Children's comprehension of scopally ambiguous questions. In *Proceedings of the 34rd West Coast Conference on Formal Linguistics*, (pp. 177–186). Cascadilla Press.

Dickson, N., Futrell, R., & Pearl, L. (2024). I Forgot but It's Okay: Learning about Island Constraints under Child-Like Memory Constraints. In AlThagafi, H. A. A. & Ray, J. (Eds.), *Proceedings of the 48th annual Boston University Conference on Language Development*, (pp. 169–183)., Somerville, MA. Cascadilla Press.

Dickson, N., Pearl, L., & Futrell, R. (2022). Learning constraints on *wh*-dependencies by learning how to efficiently represent *wh*-dependencies: A developmental modeling investigation with Fragment Grammars. *Proceedings of the Society for Computation in Linguistics*, *5*(1), 220–224.

Feldman, N., Griffiths, T., Goldwater, S., & Morgan, J. (2013). A role for the developing lexicon in phonetic category acquisition. *Psychological Review*, *120*(4), 751–778.

Fitz, H. & Chang, F. (2017). Meaningful questions: The acquisition of auxiliary inversion in a connectionist model of sentence production. *Cognition*, *166*, 225–250.

Fodor, J. D. (1998a). Parsing to Learn. *Journal of Psycholinguistic Research*, *27*(3), 339–374.

Fodor, J. D. (1998b). Unambiguous Triggers. *Linguistic Inquiry*, *29*, 1–36.

Fodor, J. D. (2017). Ambiguity, parsing, and the evaluation measure. *Language Acquisition*, *24*(2), 85–99.

Fodor, J. D. & Sakas, W. G. (2005). The subset principle in syntax: Costs of compliance. *Journal of Linguistics*, *41*(03), 513–569.

Fodor, J. D. & Sakas, W. G. (2017). Learnability. In I. Roberts (Ed.), *The Oxford Handbook of Universal Grammar* (pp. 249–269). Oxford, UK: Oxford University.

Fodor, J. D., Sakas, W. G., & Hoskey, A. (2007). Implementing the subset principle in syntax acquisition: Lattice-based models. In *Proceedings of the Second European Cognitive Science Conference*, (pp. 161–166)., Hove, UK. Lawrence Erlbaum.

Foraker, S., Regier, T., Khetarpal, N., Perfors, A., & Tenenbaum, J. (2009). Indirect Evidence and the Poverty of the Stimulus: The Case of Anaphoric One. *Cognitive Science*, *33*, 287–300.

Forsythe, H. & Pearl, L. (2020). Immature representation or immature deployment? Modeling child pronoun resolution. In *Proceedings of the Society for Computation in Linguistics 3*, (pp. 22–26).

Frank, R., Mathis, D., & Badecker, W. (2013). The acquisition of anaphora by simple recurrent networks. *Language Acquisition*, *20*(3), 181–227.

Freudenthal, D., Gobet, F., & Pine, J. M. (2024). MOSAIC+: A Crosslinguistic Model of Verb-Marking Errors in Typically Developing Children and Children With Developmental Language Disorder. *Language Learning*, *74*(1), 111–145.

Freudenthal, D., Pine, J., & Gobet, F. (2010). Explaining quantitative variation in the rate of Optional Infinitive errors across languages: a comparison of MOSAIC and the Variational Learning Model. *Journal of Child Language*, *37*(03), 643–669.

Freudenthal, D., Pine, J. M., Aguado-Orea, J., & Gobet, F. (2007). Modeling the developmental patterning of finiteness marking in English, Dutch, German, and Spanish using MOSAIC. *Cognitive Science*, *31*(2), 311–341.

Freudenthal, D., Pine, J. M., & Gobet, F. (2009). Simulating the referential properties of Dutch, German, and English root infinitives in MOSAIC. *Language Learning and Development*, *5*(1), 1–29.

Freudenthal, D., Pine, J. M., Jones, G., & Gobet, F. (2015). Defaulting effects contribute to the simulation of cross-linguistic differences in Optional Infinitive errors. In *Proceedings of the 37th annual meeting of the Cognitive Science Society*, (pp. 746–751)., Pasadena.

Futrell, R. & Mahowald, K. (2025). How Linguistics Learned to Stop Worrying and Love the Language Models. *https://doi.org/10.48550/arXiv.2501.17047*.

Futrell, R., Mahowald, K., & Gibson, E. (2015). Quantifying word order freedom in dependency corpora. In *Proceedings of the third international conference on dependency linguistics (DepLing 2015)*, (pp. 91–100).

Gibson, E. & Wexler, K. (1994). Triggers. *Linguistic Inquiry*, *25*(4), 407–454.

Gold, M. E. (1967). Language identification in the limit. *Information and Control*, *10*(5), 447–474.

Gualmini, A. (2004). Some knowledge children don't lack. *Linguistics*, *42*(5), 957–982.

Gualmini, A., Hulsey, S., Hacquard, V., & Fox, D. (2008). The question–answer requirement for scope assignment. *Natural Language Semantics*, *16*, 205–237.

Hahn, M. & Rofin, M. (2024). Why are sensitive functions hard for Transformers? In Ku, L., Martins, A., & Srikumar, V. (Eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, (pp. 14973–15008)., Bangkok, Thailand. Association for Computational Linguistics.

Heinz, J. (2009). On the role of locality in learning stress patterns. *Phonology*, *26*(2), 303–351.

Heinz, J. (2016). Computational Theories of Learning and Developmental Psycholinguistics. In J. Lidz, W. Synder, & J. Pater (Eds.), *The Oxford Handbook of Developmental Linguistics* (pp. 633–663). Oxford University Press.

Heinz, J. & Rawski, J. (2024). History of phonology: Learnability. In E. Dresher & H. ven der Hulst (Eds.), *The Oxford Handbook of the History of Phonology*. Oxford University Press.

Heinz, J. N. (2007). *Inductive learning of phonotactic patterns*. PhD thesis, University of California, Los Angeles.

Horning, J. J. (1969). A study of grammatical inference. Technical report, Stanford University, Department of Computer Science.

Hsu, A. S. & Chater, N. (2010). The logical problem of language acquisition: A probabilistic perspective. *Cognitive Science*, *34*(6), 972–1016.

Hsu, A. S., Chater, N., & Vitányi, P. (2013). Language Learning From Positive Evidence, Reconsidered: A Simplicity-Based Approach. *Topics in Cognitive Science*, *5*(1), 35–55.

Hsu, A. S., Chater, N., & Vitányi, P. M. (2011). The probabilistic analysis of language acquisition: Theoretical, computational, and experimental analysis. *Cognition*, *120*(3), 380–390.

Huang, C.-T. J. (1982). *Logical relations in Chinese and the theory of grammar*. PhD thesis, MIT, Cambridge, MA.

Kallini, J., Papadimitriou, I., Futrell, R., Mahowald, K., & Potts, C. (2024). Mission: Impossible language models. In Ku, L.-W., Martins, A., & Srikumar, V. (Eds.), *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, (pp. 14691–14714)., Bangkok, Thailand. Association for Computational Linguistics.

Koplenig, A., Meyer, P., Wolfer, S., & Müller-Spitzer, C. (2017). The statistical trade-off between word order and word structure–Large-scale evidence for the principle of least effort. *PloS One*, *12*(3), e0173614.

Lan, N., Chemla, E., & Katzir, R. (2024). Large language models and the argument from the poverty of the stimulus. *Linguistic Inquiry*, 1–56.

Lasnik, H. & Saito, M. (1984). On the nature of proper government. *Linguistic Inquiry*, *15*, 235–289.

Legate, J. & Yang, C. (2007). Morphosyntactic Learning and the Development of Tense. *Linguistic Acquisition*, *14*(3), 315–344.

Mahowald, K., Diachek, E., Gibson, E., Fedorenko, E., & Futrell, R. (2023). Grammatical cues to subjecthood are redundant in a majority of simple clauses across languages. *Cognition*, *241*.

Maitra, S. & Perkins, L. (2023). Filtering input for learning constrained grammatical variability: The case of Spanish word order. In *Proceedings of the Society for Computation in Linguistics 2023*, (pp. 108–120).

Marr, D. (1982). *Vision*. San Francisco, CA: W.H. Freeman.

Musolino, J. & Lidz, J. (2006). Why children aren't universally successful with quantification. *Linguistics*, *44*(4), 817–852.

Nakkiran, P., Kaplun, G., Bansal, Y., Yang, T., Barak, B., & Sutskever, I. (2021). Deep double descent: Where bigger models and more data hurt. *Journal of Statistical Mechanics: Theory and Experiment*, *2021*(12), 124003.

Nguyen, E. & Pearl, L. (2018). Do You Really Mean It? Linking Lexical Semantic Profiles and the Age of Acquisition for the English Passive. In *Proceedings of the 35th West Coast Conference on Formal Linguistics*, (pp. 288–295–xxx)., Somerville, MA. Cascadilla Press.

Nguyen, E. & Pearl, L. (2019). Using Developmental Modeling to Specify Learning and Representation of the Passive in English Children. In Brown, M. M. & Dailey, B. (Eds.), *Proceedings of the Boston University Conference on Language Development 43*, (pp. 469–482)., Somerville, MA. Cascadilla Press.

Nguyen, E. & Pearl, L. (2021). The link between lexical semantic features and children's comprehension of English verbal *be*-passives. *Language Acquisition*, *28*(4), 433–450.

Niyogi, P. (2006). *The computational nature of language learning and evolution*. MIT press Cambridge, MA.

Niyogi, P. & Berwick, R. C. (1996). A language learning model or finite parameter spaces. *Cognition*, *61*, 161–193.

Omaki, A., Davidson White, I., Goro, T., Lidz, J., & Phillips, C. (2014). No fear of commitment: Children's incremental interpretation in English and Japanese wh-questions. *Language Learning and Development*, *10*(3), 206–233.

Orita, N., McKeown, R., Feldman, N. H., Lidz, J., & Boyd-Graber, J. (2013). Discovering pronoun categories using discourse information. In *Proceedings of the 35th annual conference of the cognitive science society*.

Osherson, D., de Jongh, D., Martin, E., & Weinstein, S. (1997). Formal Learning Theory. In *Handbook of Logic and Language* (pp. 737–775). Elsevier.

Osherson, D. N., Stob, M., & Weinstein, S. (1983). Formal Theories of Language Acquisition: Practical and Theoretical Perspectives. In *IJCAI*, (pp. 566–572).

Osherson, D. N., Stob, M., & Weinstein, S. (1986). *Systems that learn: An introduction to learning theory for cognitive and computer scientists.* Cambridge, MA: The MIT Press.

Papafragou, A. & Musolino, J. (2003). Scalar implicatures: experiments at the semantics–pragmatics interface. *Cognition*, *86*(3), 253–282.

Pearl, L. (2007). *Necessary Bias in Natural Language Learning*. PhD thesis, University of Maryland, College Park, College Park, MD.

Pearl, L. (2009). Learning English Metrical Phonology: When Probability Distributions Are Not Enough. In J. Crawford, K. Otaki, & M. Takahashi (Eds.), *Proceedings of the 3rd Conference on Generative Approaches to Language Acquisition, North America (GALANA 2008)* (pp. 200–211). Somerville, MA: Cascadilla Press.

Pearl, L. (2011). When unbiased probabilistic learning is not enough: Acquiring a parametric system of metrical phonology. *Language Acquisition*, *18*(2), 87–120.

Pearl, L. (2014). Evaluating learning strategy components: Being fair. *Language*, *90*(3), e107–e114.

Pearl, L. (2017). Evaluation, use, and refinement of knowledge representations through acquisition modeling. *Language Acquisition*, *24*(2), 126–147.

Pearl, L. (2021). Theory and predictions for the development of morphology and syntax: A Universal Grammar+ statistics approach. *Journal of Child Language*, *48*(5), 907–936.

Pearl, L. (2022). Poverty of the Stimulus Without Tears. *Language Learning and Development*, *18*(4), 415–454.

Pearl, L. (2023a). Computational cognitive modeling for syntactic acquisition: Approaches that integrate information from multiple places. *Journal of Child Language*, *50*(6), 1353–1373.

Pearl, L. (2023b). Modeling syntactic acquisition. In Jon Sprouse (Ed.), *Oxford Handbook of Experimental Syntax* (pp. 209–270). Oxford.

Pearl, L. & Bates, A. (2022). A new way to identify if variation in children's input could be developmentally meaningful: Using computational cognitive modeling to assess input across socio-economic status for syntactic islands. *Journal of Child Language*, *51*(4), 800–833.

Pearl, L. & Forsythe, H. (2022). Inaccurate representations, inaccurate deployment, or both? Using computational cognitive modeling to investigate the development of pronoun interpretation in Spanish. *lingbuzz*.

Pearl, L. & Lidz, J. (2009). When domain-general learning fails and when it succeeds: Identifying the contribution of domain-specificity. *Language Learning and Development*, *5*(4), 235–265.

Pearl, L., Lu, K., & Haghighi, A. (2017). The character in the letter: Epistolary attribution in Samuel Richardson's *Clarissa*. *Digital Scholarship in the Humanities*, *32*(2), 355–376.

Pearl, L. & Mis, B. (2011). How Far Can Indirect Evidence Take Us? Anaphoric One Revisited. In Carlson, L., Höschler, C., & Shipley, T. (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, (pp. 879–884)., Austin, TX. Cognitive Science Society.

Pearl, L. & Mis, B. (2016). The role of indirect positive evidence in syntactic acquisition: A look at anaphoric one. *Language*, *92*(1), 1–30.

Pearl, L. & Sprouse, J. (2013a). Computational Models of Acquisition for Islands. In J. Sprouse & N. Hornstein (Eds.), *Experimental Syntax and Islands Effects* (pp. 109–131). Cambridge: Cambridge University Press.

Pearl, L. & Sprouse, J. (2013b). Syntactic islands and learning biases: Combining experimental syntax and computational modeling to investigate the language acquisition problem. *Language Acquisition*, *20*, 19–64.

Pearl, L. & Sprouse, J. (2015). Computational modeling for language acquisition: A tutorial with syntactic islands. *Journal of Speech, Language, and Hearing Research*, *58*, 740–753.

Pearl, L. & Sprouse, J. (2019). Comparing solutions to the linking problem using an integrated quantitative framework of language acquisition. *Language*, *95*(4), 583–611.

Pearl, L. & Sprouse, J. (2021). The acquisition of linking theories: A Tolerance and Sufficiency Principle approach to deriving UTAH and rUTAH. *Language Acquisition*, *28*(3), 294–325.

Pearl, L. & Weinberg, A. (2007). Input filtering in syntactic acquisition: Answers from language change modeling. *Language Learning and Development*, *3*(1), 43–72.

Perfors, A., Tenenbaum, J., & Regier, T. (2011). The learnability of abstract syntactic principles. *Cognition*, *118*, 306–338.

Perkins, L., Feldman, N., & Lidz, J. (2024). Mind the Gap: Learning the Surface Forms of Movement Dependencies. lingbuzz/008171.

Perkins, L., Feldman, N. H., & Lidz, J. (2022). The Power of Ignoring: Filtering Input for Argument Structure Acquisition. *Cognitive Science*, *46*(1), 1–43.

Perkins, L. & Hunter, T. (2023). Noise-tolerant learning as selection among deterministic grammatical hypotheses. In *Proceedings of the Society for Computation in Linguistics 2023*, (pp. 186–198).

Pijpops, D. & Zehentner, E. (2022). How redundant is language really? Agent-recipient disambiguation across time and space. *Glossa: A Journal of General Linguistics*, *7*(1), 1–41.

Portelance, E. & Jasbi, M. (2024). The roles of neural networks in language acquisition. *Language and Linguistics Compass*, *18*(6), e70001.

Power, A., Burda, Y., Edwards, H., Babuschkin, I., & Misra, V. (2022). Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*.

Pungor, J. R., Allen, V. A., Songco-Casey, J. O., & Niell, C. M. (2023). Functional organization of visual responses in the octopus optic lobe. *Current Biology*, *33*(13), 2784–2793.

Rajalingham, R., Issa, E. B., Bashivan, P., Kar, K., Schmidt, K., & DiCarlo, J. J. (2018). Large-scale, high-resolution comparison of the core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *Journal of Neuroscience*, *38*(33), 7255–7269.

Rawski, J. (2021). *Structure and Learning in Natural Language*. PhD thesis, State University of New York at Stony Brook.

Reali, F. & Christiansen, M. (2005). Uncovering the Richness of the Stimulus: Structure Dependence and Indirect Statistical Evidence. *Cognitive Science*, *29*, 1007–1028.

Regier, T. & Gahl, S. (2004). Learning the unlearnable: The role of missing evidence. *Cognition*, *93*, 147–155.

Rizzi, L. (1982). Violations of the wh-island constraint and the subjacency condition. In L. Rizzi (Ed.), *Issues in Italian Syntax*. Dordrecht, NL: Foris.

Ross, J. (1967). *Constraints on variables in syntax*. PhD thesis, MIT, Cambridge, MA.

Sakas, W. (2016). Computational approaches to parameter setting in generative linguistics. In J. Lidz, W. Snyder, & J. Pater (Eds.), *The Oxford Handbook of Developmental Linguistics* (pp. 696–724). Oxford, UK: Oxford University Press.

Sakas, W. & Fodor, J. (2012). Disambiguating Syntactic Triggers. *Language Acquisition*, *19*(2), 83–143.

Sakas, W. & Fodor, J. D. (2001). The Structural Triggers Learner. In S. Bertolo (Ed.), *Language Acquisition and Learnability* (pp. 172–233). Cambridge, UK: Cambridge University Press.

Sakas, W. & Nishimoto, E. (2002). Search, Structure or Statistics? A Comparative Study of Memoryless Heuristics for Syntax Acquisition. Manuscript.

Savinelli, K., Scontras, G., & Pearl, L. (2017). Modeling scope ambiguity resolution as pragmatic inference: Formalizing differences in child and adult behavior. In *Proceedings of the 39th annual meeting of the Cognitive Science Society*, London, UK. Cognitive Science Society.

Savinelli, K., Scontras, G., & Pearl, L. (2018). Exactly two things to learn from modeling scope ambiguity resolution: Developmental continuity and numeral semantics. In *Proceedings of the 8th workshop on Cognitive Modeling and Computational Linguistics (CMCL 2018)*, (pp. 67–75).

Scontras, G. & Pearl, L. S. (2021). When pragmatics matters more for truth-value judgments: An investigation of quantifier scope ambiguity. *Glossa: A journal of General Linguistics*, *6*(1).

Tîrnăucuă, C. (2008). A note on the relationship between different types of correction queries. In *Grammatical Inference: Algorithms and Applications: 9th International Colloquium, ICGI 2008 Saint-Malo, France, September 22-24, 2008 Proceedings 9*, (pp. 213–223). Springer.

Torrego, E. (1984). On inversion in Spanish and some of its effects. *Linguistic Inquiry*, *15*, 103–129.

Valiant, L. (1984). A theory of the learnable. *Communications of the ACM*, 1134–1142.

Wexler, K. & Culicover, P. (1980). *Formal principles of language acquisition*. Cambridge, MA: MIT Press.

Wharton, R. (1974). Approximate language identification. *Information and Control*, *26*(3), 236–255.

Wiehagen, R. (1977). Identification of formal languages. In *International Symposium on Mathematical Foundations of Computer Science*, (pp. 571–579). Springer.

Wilcox, E. G., Futrell, R., & Levy, R. (2023). Using computational models to test syntactic learnability. *Linguistic Inquiry*, 1–44.

Yang, C. (2002). *Knowledge and Learning in Natural Language*. Oxford, UK: Oxford University Press.

Yang, C. (2004). Universal grammar, statistics or both? *Trends in Cognitive Science*, *8*(10), 451–456.

Yang, C. (2012). Computational models of syntactic acquisition. *WIREs Cognitive Science*, *3*, 205–213.

Yang, X., Aoyama, T., Yao, Y., & Wilcox, E. (2025). Anything Goes? A Crosslinguistic Study of (Im)possible Language Learning in LMs. *arXiv preprint arXiv:2502.18795*.

Yedetore, A., Linzen, T., Frank, R., & McCoy, R. T. (2023). How poor is the stimulus? Evaluating hierarchical generalization in neural networks trained on child-directed speech. In Rogers, A., Boyd-Graber, J., , & Okazaki, N. (Eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, (pp. 9370–9393)., Toronto, Canada. Association for Computational Linguistics.