Modeling the Relational Shift in Melodic Processing of Young Children

Ahnate Lim (ahnate@hawaii.edu)
Department of Psychology, University of Hawaii at Manoa
2530 Dole Street, Honolulu, HI 96822 USA

Leonidas A. A. Doumas (leonidas@hawaii.edu)
Department of Psychology, University of Hawaii at Manoa
2530 Dole Street, Honolulu, HI 96822 USA

Scott Sinnett (ssinnett@hawaii.edu)
Department of Psychology, University of Hawaii at Manoa
2530 Dole Street, Honolulu, HI 96822 USA

Abstract

As a ubiquitous trend in the cognitive development of children, the 'relational shift' accounts for a change in preference for absolute percepts towards a preference for relational percepts, and is observed across a wide variety of domains. Extensive evidence indicates that this prepotency for relational processing is also observed in how children process melodies. When recalling melodies, younger children typically recall more absolute pitch properties than older children, while the exact opposite occurs in older children. Using DORA (*Discovery Of Relations by Analogy*; Doumas et al., 2008), a domain-general symbolic-connectionist model of relation learning, we simulated the relational shift in melodic perception of children age 3-6 years based on an experiment by Sergeant and Roche (1973). DORA's performance matched the children's well, suggesting common developmental and perceptual mechanisms between the relational shift in melodic processing and the shift seen across other domains.

Keywords: Melodic perception; relation learning; development; relational shift; absolute pitch; computational modeling; DORA.

Introduction

One of the fundamental cross-domain trends in human development is characterized by a qualitative transformation, or shift, in how children process information. Evidence from developmental psychology overwhelmingly indicates that while children initially attend to, recall, and reason about absolute perceptual properties, around the age of 4-6 they begin to rely on structured relational properties (Allport, 1924; Gentner & Rattermann, 1991; Halford, 2005; Pollack, 1969; Vernon, 1940). This shift has been observed in areas such as language (Gentner, 1988), spatial tasks (Case & Khanna, 1981; DeLoache, Sugarman, & Brown, 1985), number comprehension (Gelman & Gallistel, 1978; Michie, 1985), and visual shape perception (Abecassis, Sera, Yonas, & Schwade, 2001), to name but a few. This phenomenon has been termed the 'relational shift,' as the characteristic trend is towards greater reliance on relational attributes as children mature.

Consistent with the developmental trajectory for the relational shift in other domains, in the domain of music children also develop from initially processing more absolute aspects of melodies to processing more relational aspects as they grow older. In an especially revealing study, Sergeant and Roche (1973) trained three groups of children from the age of three to six to reproduce melodies from invariant recordings. When the children were required to recall the melodies one week later, the younger group reproduced the absolute pitches more accurately than the older group, while the older group reproduced the relational aspects (melodic shape, interval sizes, and tonality) more accurately than the younger group. This perceptual shift and exchange in proficiency levels between the recall of absolute and relational musical aspects in younger and older children has been replicated in many other studies as well (Saffran, 2003; Saffran & Griepentrog, 2001; Sergeant, 1969; Sergeant & Roche, 1973; Stalinski & Schellenberg, 2010; Takeuchi & Hulse, 1993).

Given the prevalence of the relational shift across multiple domains, it is reasonable to assert that any comprehensive theory or model of cognitive development must necessarily account for this phenomenon. One of the models of higher cognition that has successfully been used to account for the relational shift in development is DORA (Discovery Of Relations by Analogy; Doumas, Hummel, & Sandhofer, 2008). DORA has been used to simulate the relational shift in visual shape perception (Doumas & Hummel, 2010), analogical problems (Doumas, Morrison, & Richland, 2009; Morrison, Doumas, & Richland, 2011), categorical reasoning, spatial reasoning, general relational reasoning, and progressive alignment (Doumas et al., 2008).

In this study, we aim to understand how the relational shift in melodic processing occurs in children. We hypothesize that the same processes that cause the relational shift in other domains are also responsible for the shift in the domain of melodic processing. Specifically, we propose that as children learn about the world, they increasing rely on relational invariants in the environment. This reliance is itself a direct result of the cognitive processes that allow for relational invariants to be detected in the first place. That is, equipped with a cognitive architecture that allows for intersection discovery of shared properties, the natural trend over time (i.e., repeated exposure) is to preferentially perceive the world in terms of these regularly occurring relational invariants (Doumas & Hummel, 2010; Doumas et al., 2008).

This raises the question as to how relational invariants are discovered in the first place. Our proposal for this mechanism of discovery is instantiated in DORA's symbolic-connectionist architecture, and has been used to account for how melodic perception occurs in infants (Lim, Doumas, & Sinnett, 2012). Consequently, providing an account for the relational shift in melodic processing may also help to shed light on other issues. For instance, the argument could be made against the existence of a musical relational shift by citing evidence of infants' ability to detect relational properties from melodies (Plantinga & Trainor, 2005; Stalinski & Schellenberg, 2012; Trehub, Bull, & Thorpe, 1984). That is, given that the relational shift indicates that younger children preferentially process melodies based on absolute percepts (i.e., absolute pitch), would evidence of infants ability to process melodies based on relative percepts (i.e., relative pitch) not be contradictory? Since DORA has been used to simulate the latter phenomenon (Lim et al., 2012), by using DORA to simulate the former phenomenon (i.e., the relational shift in musical processing), we hope to provide an answer to this question as well.[1]

The LISA/DORA models

LISA *(Learning and Inference with Schemas and Analogies*; Hummel & Holyoak, 1997, 2003) is a symbolic-connectionist model of analogy and relational reasoning. DORA is an extension of LISA that learns structured (i.e., symbolic) representations of relations from unstructured inputs. That is, DORA provides an account of how the structured relational representations LISA uses to perform relational reasoning can be learned from examples. At present, DORA accounts for over 30 phenomena from the literature on relational learning, and cognitive development, and as it learns representations of relations it develops into LISA and can simulate the additional 40+ phenomena in relational thinking for which LISA accounts for (e.g., Doumas et al., 2008). In the following, we provide a very brief description of the LISA/DORA models (for full details, see Hummel & Holyoak, 1997, 2003; Doumas et al., 2008)

LISAese Representations In LISA (and DORA *after* it has gone through learning), relational structures are represented by a hierarchy of distributed and localist codes (see Figure 1). At the bottom, "semantic" units (small circles in Figure 1) represent the features of objects and roles in a distributed fashion. At the next level, these distributed representations are connected to localist units (POs) representing individual predicates (or roles) and objects (triangles and larger circles in Figure 1). Localist role-binding units (RBs; rectangles in Figure 1) link object and role units into specific role-filler bindings. At the top of the hierarchy, localist P units (ovals in Figure 1) link RBs into whole relational propositions.

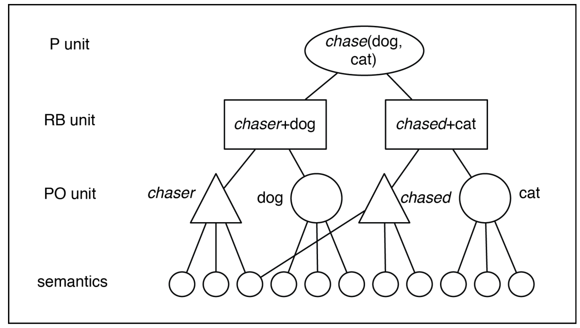

Figure 1: LISA/DORA representation of the proposition, *chase* (dog, cat).

Relational structures (or propositions) are divided into two mutually exclusive sets: a driver and recipient(s). In LISA/DORA, the sequence of firing events is controlled by the driver. Specifically, one (or at most three) proposition(s) in the driver become(s) active (i.e., enter working memory). When a proposition enters working memory, role-filler bindings must be represented dynamically on the units that maintain role-filler independence (i.e., POs and semantic units) to allow for reusability of units and preservation of similarity across different bindings (Hummel & Holyoak, 1997). In LISA, binding information is carried by synchrony of firing (with roles firing simultaneously with their fillers). In DORA, binding information is carried by systematic asynchrony of firing, with bound role-filler pairs firing in direct sequence (for details, see Doumas et al., 2008).[2]

Relational Learning In broadest strokes, DORA learns structured representations by comparing objects to isolate their shared properties and to represent these shared properties as explicit structures. More specifically, DORA starts with simple feature-vector representations of objects (i.e., a node connected to set of features describing that object; large and small circles from Figure 1). When DORA compares one object to another, corresponding elements (i.e., shared features) of the two representations fire simultaneously (Figure 2a). Any semantic features common to both objects receive twice as much input and thus become roughly twice as active as features connected to one but not the other (Figure 2b). By recruiting a new PO unit and learning connections between that unit and active semantics via Hebbian learning (wherein the strength of connections is a function of the units' activation), DORA learns stronger connections between the new PO unit and more active

[1] Due to spatial constraints, we provide only summary information on melodic and relational processing here, for more background on melodic processing, including details about absolute and relative pitch and the other features used within these simulations, see Lim et al. (2012).

[2] Asynchrony-based binding allows role and filler to be coded by the same pool of semantic units, which allows DORA to learn representations of relations from representations of objects (Doumas et al., 2008).

semantic units (Figure 2c). The new PO thus becomes an explicit representation of the featural overlap of the compared objects and can act as a single place predicate, taking other object representations as arguments to form role-filler pairs (Figure 2d; see also Doumas et al., 2008). Applied iteratively, this process allows DORA to learn structured explicit single-place predicate representations of any properties that compared objects may share. Comparison also allows DORA to learn representations of multi-place relations by linking sets of constituent role-filler pairs into relational structures (i.e., to learn the *chases* relation by linking together representations of the roles *chaser* and *chased*; see Doumas et al., 2008 for details). Moreover, when DORA has learned representations of whole relational structures, this intersection discovery algorithm allows it to learn schemas by comparing sets of structural relations to one another. For example, if after DORA has learned about a dog chasing a cat (*chase* (dog, cat)) and a boy chasing a girl (*chase* (boy, girl)), it can compare these and learn a representation coding for the intersection of the two *chase* relations and their arguments, or *chase* (generic-object1, generic-object2).

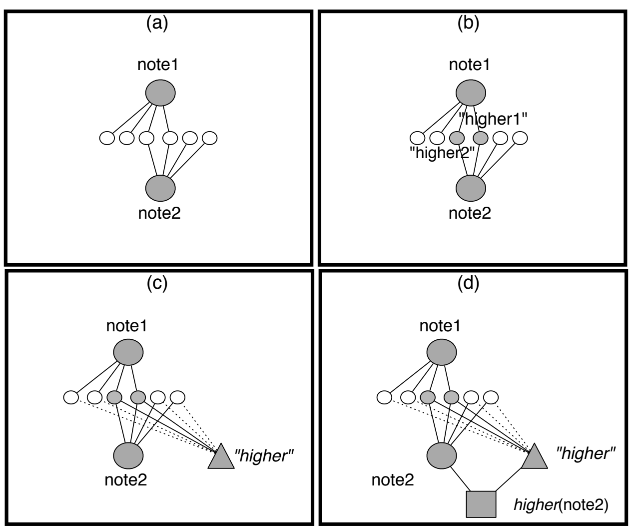

Figure 2: The process through which DORA learns a single-place predicate representation of “higher” from two musical notes.

Mapping For the purpose of analogical mapping, LISA/DORA learns *mapping connections* between units coactive of the same type in the driver and recipient (e.g., between PO units in the driver and PO units in the recipient). These connections grow whenever corresponding units in the driver and recipient are active simultaneously. They permit LISA to learn the correspondences between matching structures in separate analogs. They also permit correspondences learned early in mapping to influence the correspondences learned later.

Methods

In this section we describe the Sergeant and Roche (1973) study, followed by the details and outcomes of DORA’s simulations.

Task Description

In Sergeant and Roche’s (1973) cross-sectional study, children were divided into three groups: one group with children between 3 to 4 years, one with children of 5 years, and one with children of 6 years. All groups received the same training and testing procedures. They were trained to vocally reproduce three melodies from an invariant recording in six training sessions spread out over three weeks. All children were given the exact same melodies at each training session. Each melody lasted for 8 or 16 bars.

One week after training, the children were then asked to vocally recall the melodies, which were tape recorded and scored by two independent judges on perceptual dimension (pitch accuracy), and conceptual dimensions (melodic shape, intervals, and tonality).

Simulations

In the first simulation, we simulated the development of representations of individual relations that could define auditory sequence from experience with the world. In the second simulation, we used the representations DORA learned during the first simulation to simulate the behavior of Sergeant and Roche’s (1973) subjects. Crucially, these two simulations were interleaved, which allowed us to test

Simulation Part 1 In the first simulation we tested DORA’s ability to learn relational concepts from examples. This simulation proceeded like several simulations of relation learning from our previous work (e.g., Doumas & Hummel, 2005, 2010; Doumas et al., 2008; Doumas et al., 2009). We started DORA with representations of 100 objects (represented as PO units) attached to random sets of features (chosen from a pool of 100). We then defined 4 relations (those that could be used to describe a melodic sequence, e.g., contour (*higher/lower*), and interval (*long-interval*, *short-interval, medium-interval*)).

Each relation transformation consisted of two roles each with three semantic features (e.g., for the *higher* relation, both the roles *above* and *below* were each defined by three specific semantic units). Each of the 100 objects was attached to the features of between 2 and 4 relational roles chosen at random such that if an object was part of a relation, it was attached to the features of one of the roles, chosen at random. For example, object1 might be attached to the features for *above* (one role of the relation *higher*) and *start-long-interval* (the agent role of the relation *long-interval*). We presented DORA with sets of objects selected at random, and allowed it to compare the objects and learn from the results (as per DORA’s relation learning algorithm). As DORA learned new representations it would also use these representations to make subsequent comparisons. For instance, if DORA learned an explicit representation of the property *above* by comparing two objects both attached to the features of *above*, it could use this new representation for future comparisons. On each trial we selected between 2 and 6 representations and let DORA compare them and learn from the results (i.e., perform predication, and relation learning routines). We assume that this act of inspection and comparison is similar to what happens when children encounter objects in the world—where objects are part of several relations—and learns from these experiences.

We ran 600 learning trials and measured the quality of the representations DORA had learned after each 100 trials. Quality was calculated as the mean of connection weights to relevant features (i.e., those defining a specific transformation or role of a transformation) divided by the mean of all other connection weights + 1 (1 was added to the mean of all other connection weights to normalize the quality measure to between 0 and 1). A higher quality denoted stronger connections to the semantics defining a specific transformation relative to all other connections (i.e., a more pristine representation of the transformation). Figure 3 indicates the quality of the representations DORA learned at each level of iteration. Early in learning, DORA's representations are 'dirty' in that it's representations of relations and their roles are also highly connected to extraneous features specific to the instances from which the representations are learned. These representations are consequently very context dependent. As learning progresses however, DORA's representations become progressively more refined. By the end of learning, DORA has learned representations of relations and their roles that are context-independent, connected strongly to only the features specific to the particular relational roles defining the relation and very weakly connected to context features. Thus, in time DORA can use these representations to reason about instances regardless of context, like older children and adults (see, e.g., Doumas et al., 2008).

For the analysis herein, the 'quality' of DORA's representations (how relationally clean or context dependent they are) is considered an analogous measurement to the vocal reproductions of the children in Sergeant and Roche's (1973) study. That is, more pristine representations in DORA would be analogous to children reproducing melodies with more conceptual (relational) dimensions, whereas dirty presentations in DORA would be analogous to children reproducing melodies with more perceptual (absolute) dimensions.

Simulation Part 2 During the second simulation we simulated Sergeant and Roche's (1973) training and test conditions. We created a 20 note melody represented as 20 PO units attached to features indicating absolute frequency (between f1 and f24), the note's place in the sequence (1-20), two semantics describing whether the note is higher (*above*) or lower (*below*) the previous note in the sequence, two semantics describing the relative interval from the previous note (*high-, medium-, low-interval*), a semantic describing the absolute interval from the previous note, and four random features (from a pool of 100). The features represent the properties that infants, children, and adults are capable of representing about melody (Thorpe & Trehub, 1989; Trehub et al., 1984). Importantly, all of the frequency direction (higher/lower) and frequency interval, both absolute and relative, can be generated from raw frequency values (i.e. sensory input) using a simple comparator circuit described in Doumas et al. (2008) and Hummel and Biederman (1992).

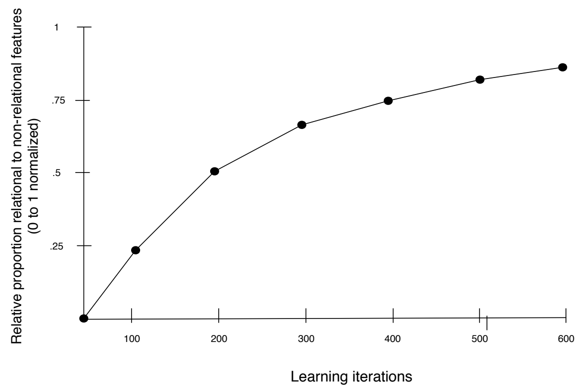

Figure 3: The quality of DORA's representations as a function of learning iterations.

During training, we presented DORA with the note sequence and allowed it to fire each two note sequence in the melody (e.g., notes 1 and 2, then notes 2 and 3). During each two note firing sequence DORA attempted to retrieve relations from LTM describing the sequence (these representations were the same as those DORA had learned during part one of the simulation; see below for details). If DORA successfully retrieved a relation from LTM, DORA predicated the respective roles of the relation about the notes in the sequence. For example, if a two note sequence caused DORA to recall the *higher* (*x*, *y*) relation from LTM (consisting of the roles *above*(*x*) and *below*(*y*)), DORA would link the *above* PO to the note that was higher with an RB unit, and the *below* PO to the note that was lower. This process reflects our assumption that children and adults attempt to understand melodies using representations at their disposal. After DORA has attempted to classify the 2 note sequences in the melody, DORA stores the resulting representation in LTM.

Importantly, to simulate 4, 5, and 6-year olds, we used representations that DORA had learned during the first part of the simulation in DORA's LTM. Specifically, to simulate the representations of 4 year olds, we used the representations that DORA had learned after 200 training iterations, to simulate 5 year olds we used the representations DORA had learned after 400 iterations, and to simulate 6 year-olds we used the representations DORA had learned after 600 iterations. At each age we also included distractor predicates describing extraneous properties (e.g., loudness, timbre, etc.) in LTM. For every relevant relation in DORA's LTM (i.e., relations describing *higher* and relative interval) we also included 2 irrelevant relations. Our addition of distractor relations in LTM instantiates our assumption that children learn about multiple relations at the same time during development.

We trained DORA in this manner six times (reflecting the six training sessions from the Sergeant & Roche (1973) study). After the second training session, and after each subsequent training session, DORA compared the representation it had learned during training to the representation it had learned during the previous training session and learned a new representation (or a schema) using it's learning algorithm.

To simulate the testing phase from Sergeant and Roche's (1973) study, we examined the representation of the melody DORA has learned after the six training sessions. Four-year-old DORA's relational representations were quite dirty and tied to the semantics of the objects from which they had been learned. DORA, consequently, had difficulty retrieving these representations from memory given the melody as a context cue. As a result, the representation of the melody that DORA stores is essentially the melody itself, without much (if any) explicit relational information predicated about it. As DORA get's older (i.e., has its LTM populated with representations produces by more extensive learning during simulation part 1), it becomes more likely to retrieve and thus predicate relations about the two note sequences in the melody during training. More precisely, 4-year old DORA retrieved predicates about only 18% of the 2 note sequences it thought about, 5 year-old DORA retrieved predicates about 63% of the two note sequences it thought about, and 6 year-old DORA about 91% of the instances it though about. Importantly, the predicates in DORA's LTM that it could retrieve varied in their refinement across ages (as described above). We used the representations that DORA had learned after the six training session as a proxy for what it would recall as melody production during the test session of the Sergeant and Roche study. We evaluated these final representations for the presence of relational properties with the assumption that increases in relational properties indicate increased reliance and accuracy on the conceptual dimensions of melodic shape and relative interval. Just as the children in Sergeant and Roche (1973), early in development DORA's ratio of relational/categorical features to absolute features was low, but as DORA learned the ratio increased strongly. At age 4, the ratio value was 0.85. This value increased to 1.1 at age 5 and 1.6 at age 6. This progression very closely mirrors the change in reliance on absolute to relational properties observed in children.

Discussion

The purpose of this study was to 1) test our hypothesis that a common mechanism could potentially underly both the relational shift in melodic perception and the relational shift observed in other domains, and to 2) instantiate this mechanism within a computational model. To this end, our hypotheses were supported by DORA's simulations, which matched the behavioral data from children in Sergeant and Roche's (1973) study. To our awareness this is the first time the relational shift in melodic processing has been modeled using 1) a neurally plausible architecture, 2) a domain-general model of cognition, and 3) the first run of simulations without any parameter fittings.

Consequently, DORA's success in simulating both the relational shift in children's melodic processing in this study, and in simulating infants' ability to detect relational properties of melodies (Lim et al., 2012), provides insights into a misunderstood (what we view as nonexistent) contradiction. Specifically, the argument has been made (e.g., Stalinski & Schellenberg, 2012) that the evidence for a relational shift in melodic processing may be contradicted by findings that infants can process relational properties of melodies (for a review, see Trehub, 2001). We argue that these two findings are not contradictory, as evidence of a relational shift does not indicate that younger children cannot process relations, only that they show a *preference* for absolute pitch percepts. As they grow older this preference shifts towards relational melodic features (Takeuchi & Hulse, 1993). Our theory posits that detection of relational features in melodies by infants (and humans of all ages for that matter) is facilitated by the temporal nature of melodies (each note in the melody sequentially occurs over time), and the corresponding temporality through which our brain encodes and recalls each note (i.e., binding through asynchrony).[3]

We propose that cognitive systems (e.g., DORA) that use temporality as a binding mechanism between the individual units (notes) of a perceptual group (melodies), is inherently equipped to detect relational invariants within the group (Lim et al., 2012). Through development, learning (i.e., repeated exposure to the environment) occurs and the system inevitably detects more relational invariants, develops cleaner representations that are closer to these invariants (Simulation 1), and learns that this type of information is valuable and predictive. As a result, the system comes to prefer these types of percepts, as observed in the relational shift and predicted by DORA.

It has been proposed that the ability to detect relational properties in melodies may have a common ontogenetic origin as the ability to process vocal speech patterns, where our ability to detect relational melodic features may be a by product of our ability to detect invariants in speech (Terhardt, 1974). Additionally, our model lends support to the notion that absolute (i.e., perfect) pitch—the ability to recount specific note names or frequencies of auditory stimuli—may be a common ability in all humans that is robust in early childhood and subsequently diminishes through development, specifically as the relational shift occurs (Sergeant, 1969; for a review on aboslute pitch, see Takeuchi & Hulse, 1993). We speculate that the high correlation of musical training during childhood with

[3] Although we propose for time as a binding mechanism, we agnostically acknowledge that other mechanisms could serve a similar function. For a detailed algorithmic level account of DORA's mechanisms, see Doumas et al., 2008.

absolute pitch abilities (that subsequently endures into adulthood) may be due to increased exposure to pitch relevant stimuli as young children, and hope to examine such questions in future research.

References

Abecassis, M., Sera, M. D., Yonas, A., & Schwade, J. (2001). Whats in a shape? children represent shape variability differently than adults when naming objects. *Journal of Experimental Child Psychology, 78*(3), 213-239.

Allport, G. W. (1924). Eidetic imagery. *British Journal of Psychology. General Section, 15*(2), 99-120.

Case, R., & Khanna, F. (1981). The missing links: Stages in children's progression from sensorimotor to logical thought. *New Directions for Child and Adolescent Development, 1981*(12), 21-32.

DeLoache, J. S., Sugarman, S., & Brown, A. L. (1985). The development of error correction strategies in young children's manipulative play. *Child Development*, 928-939.

Doumas, L. A. A., & Hummel, J. E. (2005). *A symbolic-connectionist model of relation discovery.* Paper presented at the 23rd Annual Conference of the Cognitive Science Society.

Doumas, L. A. A., & Hummel, J. E. (2010). A computational account of the development of the generalization of shape information. *Cognitive science, 34*(4), 698-712.

Doumas, L. A. A., Hummel, J. E., & Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review, 115*(1), 1.

Doumas, L. A. A., Morrison, R. G., & Richland, L. E. (2009). *The Development of Analogy: Working Memory in Relational Learning and Mapping.* Paper presented at the 31st Annual Conference of the Cognitive Science Society, Amsterdam, The Netherlands.

Gelman, R., & Gallistel, C. (1978). *The child's understanding of number*. Cambridge, MA: Harvard University Press.

Gentner, D. (1988). Metaphor as structure mapping: The relational shift. *Child Development*, 47-59.

Gentner, D., & Rattermann, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on language and thought: Interrelations in development* (pp. 225-277). London: Cambridge University Press.

Halford, G. S. (2005). Development of thinking. In K. J. Holyoak & R. G. Morrison (Eds.), *The Cambridge handbook of thinking and reasoning* (pp. 529-558). New York: Cambridge University Press.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99*(3), 480.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review, 104*(3), 427.

Hummel, J. E., & Holyoak, K. J. (2003). A symbolic-connectionist theory of relational inference and generalization. *Psychological Review, 110*(2), 220.

Lim, A., Doumas, L. A. A., & Sinnett, S. (2012). *Modeling Melodic Perception as Relational Learning Using a Symbolic-Connectionist Architecture (DORA).* Paper presented at the 34th Annual Conference of the Cognitive Science Society, Sapporo, Japan.

Michie, S. (1985). Development of absolute and relative concepts of number in preschool children. *Developmental Psychology, 21*(2), 247.

Morrison, R. G., Doumas, L. A. A., & Richland, L. E. (2011). A computational account of children's analogical reasoning: balancing inhibitory control in working memory and relational representation. *Developmental Science, 14*(3), 516-529.

Plantinga, J., & Trainor, L. J. (2005). Memory for melody: Infants use a relative pitch code. *Cognition, 98*(1), 1-11.

Pollack, R. (1969). Some implications of ontogenetic changes in perception. In Elkind & Flavell (Eds.), *Studies in cognitive development* (pp. 365-407): O.U.P.

Saffran, J. R. (2003). Absolute pitch in infancy and adulthood: The role of tonal structure. *Developmental Science, 6*(1), 35-43.

Saffran, J. R., & Griepentrog, G. J. (2001). Absolute pitch in infant auditory learning: evidence for developmental reorganization. *Developmental Psychology, 37*(1), 74.

Sergeant, D. (1969). Experimental investigation of absolute pitch. *Journal of Research in Music Education, 17*(1), 135-143.

Sergeant, D., & Roche, S. (1973). Perceptual shifts in the auditory information processing of young children. *Psychology of Music, 1*(2), 39-48.

Stalinski, S. M., & Schellenberg, E. G. (2010). Shifting perceptions: Developmental changes in judgments of melodic similarity. *Developmental Psychology, 46*(6), 1799.

Stalinski, S. M., & Schellenberg, E. G. (2012). Music Cognition: A Developmental Perspective. *Topics in Cognitive Science*.

Takeuchi, A. H., & Hulse, S. H. (1993). Absolute pitch. *Psychological bulletin, 113*(2), 345.

Terhardt, E. (1974). Pitch, consonance, and harmony. *The Journal of the Acoustical Society of America, 55*, 1061.

Thorpe, L. A., & Trehub, S. E. (1989). Duration illusion and auditory grouping in infancy. *Developmental Psychology, 25*(1), 122.

Trehub, S. E., Bull, D., & Thorpe, L. A. (1984). Infants' perception of melodies: The role of melodic contour. *Child Development, 55*(3), 821-830.

Vernon, M. D. (1940). The relation of cognition and phantasy in children. *British Journal of Psychology. General Section, 31*(1), 1-21.