



Swansea University
Prifysgol Abertawe



Cronfa - Swansea University Open Access Repository

This is an author produced version of a paper published in:
European Journal of Behavior Analysis

Cronfa URL for this paper:
<http://cronfa.swan.ac.uk/Record/cronfa37398>

Paper:

Stewart, I., Slattery, B., Chambers, M. & Dymond, S. (2018). An empirical investigation of part-whole hierarchical relations. *European Journal of Behavior Analysis*, 1-20.
<http://dx.doi.org/10.1080/15021149.2017.1416525>

This item is brought to you by Swansea University. Any person downloading material is agreeing to abide by the terms of the repository licence. Copies of full text items may be used or reproduced in any format or medium, without prior permission for personal research or study, educational or non-commercial purposes only. The copyright for any work remains with the original author unless otherwise specified. The full-text must not be sold in any format or medium without the formal permission of the copyright holder.

Permission for multiple reproductions should be obtained from the original author.

Authors are personally responsible for adhering to copyright and publisher restrictions when uploading content to the repository.

<http://www.swansea.ac.uk/library/researchsupport/ris-support/>

Running head: PART-WHOLE HIERARCHICAL RELATIONS

An Empirical Investigation of Part-Whole Hierarchical Relations

Ian Stewart¹

Brian Slattery¹

Michael Chambers¹

Simon Dymond²

1 National University of Ireland, Galway

2 Swansea University, Wales

Authors' Note:

Ian Stewart, School of Psychology, National University of Ireland, Galway, Ireland;

Brian Slattery, School of Psychology, National University of Ireland, Galway,

Ireland;

Michael Chambers, School of Psychology, National University of Ireland, Galway,

Ireland;

Simon Dymond, Department of Psychology, Swansea University, Wales, UK

Corresponding Author:

Ian Stewart, School of Psychology, National University of Ireland, Galway, Ireland

Telephone: XXXXXXXXXX

E-mail: ian.stewart@nuigalway.ie

Running head: PART-WHOLE HIERARCHICAL RELATIONS

An Empirical Investigation of Part-Whole Hierarchical Relations

An Empirical Investigation of Part-Whole Hierarchical Relations

Abstract

Previous research has shown classification type (“class / member”) hierarchical responding as relational framing. The aim of this study was to use the same paradigm to investigate analysis type (“part-whole”) hierarchical responding. Ten participants were exposed to (i) a procedure involving non-arbitrarily related multi-dimensional stimuli, the aim of which was to establish two arbitrary shapes as contextual cues for ‘part of’ and ‘includes’ relational responding respectively; and (ii) a procedure that employed these cues to establish a network of arbitrary stimuli in particular hierarchical relations with each other and then test for derivation of further untrained hierarchical relations as well as for transformation of functions (TOF). Resultant patterns of relational framing consistently showed features of asymmetrical mutual entailment and transitive combinatorial entailment but, in contrast with results for classification type hierarchical responding, there was no consistent trend for TOF. Implications and future research directions are discussed.

Keywords: Hierarchical responding, relational frame theory, classification, part whole relations, transformation of functions.

An Empirical Investigation of Part-Whole Hierarchical Relations

Humans often respond in accordance with conceptual hierarchies in which higher order concepts subsume lower order ones. One example of this is hierarchical classification, in which classes of stimuli are treated as members of larger classes (e.g., ‘poodle’ is classified as a member of the category ‘dog’ while ‘dog’ is classified as a member of the category ‘animal’ etc.; see e.g., Griffiee & Dougher, 2002; Slattery, Stewart & O’Hora, 2011). Another example of a conceptual hierarchy is hierarchical part-whole analysis, in which elements are treated as parts of larger, more inclusive ‘wholes’ (e.g., ‘nail’ is categorised as part of ‘finger’ while ‘finger’ is categorised as part of ‘hand’ etc.; see e.g., Slattery & Stewart, 2014). Conceptual hierarchies are important as they allow us to organise our environment at a conceptual level so as to respond to it more effectively; indeed, such organisation can be argued to be critical to scientific thought, for example. The aim of the present study was to extend previous behavior analytic research into conceptual hierarchical responding (hereafter referred to as simply ‘hierarchical responding’).

The current research takes a Relational Frame Theory (RFT; Hayes, Barnes-Holmes & Roche, 2001; Dymond & Roche, 2013) approach to modelling hierarchical responding. From the perspective of RFT, human language and cognition can be conceptualized as learned patterns of generalized contextually controlled relational responding called relational frames. Particular frames (e.g., same, opposite, different, more / less etc.) are reinforced in the presence of particular contextual cues until eventually the cues come to reliably control their emission. Typically, a pattern of relational responding is at first primarily controlled by physical relations between stimuli (this is referred to as non-arbitrary relational responding); eventually,

however, the responding generalizes so that it is applicable to stimuli without non-arbitrary relational support. As an example, in the case of comparative (more / less) relations, a child might first learn to choose the physically larger or smaller of two objects in the presence of the auditory stimuli ‘bigger’ and ‘smaller’ respectively. Then, through exposure to multiple exemplars of this pattern in the presence of these cues, the relational response becomes abstracted such that it can be applied in conditions without an obvious formal relation; for example, after being told that ‘X1 is bigger than X2’, they will derive that ‘X2 is smaller than X1’. All frames are defined in terms of properties of mutual entailment (a relation from stimulus A to stimulus B entails a relation from B to A; e.g., in comparative framing, ‘A more than B’ entails ‘B less than A’); combinatorial entailment (the combination of relations entails further relations; e.g., in comparative framing, the combination of ‘A more than B’ and ‘B more than C’ entails ‘A more than C’ and ‘C less than A’); and transformation of (stimulus) functions (TOF; the functions of a stimulus in a relational frame may be transformed in accordance with the nature of the relation(s) derived between that stimulus and others; for example, if someone derives that A is ‘more than’ B, which has aversive functions, then A may become more aversive than B without training; see, for example, Dougher, Hamilton, Fink & Harrington, 2007).

Hierarchical responding can be conceptualized and modelled as a particular form of relational framing referred to as hierarchical relational framing. As with other frames, this might be hypothesized to originate in the training of non-arbitrary relations. For example, one such non-arbitrary relational pattern that might be important is containment. For instance, a child might learn, in one context, to describe things as being physically inside other things (e.g., ‘my hand is in my glove’) and in another, to describe things as containing other things (e.g., ‘the house contains the

doll'). This repertoire might then come under contextual control (i.e., of cues such as the words 'in' and 'contains') and generalize, developing into more abstract patterns such as classification (i.e., responding to 'members' as being contained in 'classes') or analysis (i.e., responding to 'parts' as being contained within 'wholes'). As such, one way to investigate hierarchical responding might be to establish arbitrary stimuli as cues using non-arbitrary training and then use those cues to (i) establish hierarchical relations (e.g., 'member / class') between arbitrary stimuli and (ii) gauge additional derived relations and/or properties.

To date, a number of studies have modelled hierarchical responding as hierarchical relational framing. Gil, Luciano, Ruiz and Valdivia-Salas (2012) reported a number of empirical and methodological innovations including the establishing of contextual cues for containment relations and the demonstration of a format in which responding in accordance with multiple stimulus-relations (same, different, belongs to, includes) was probed through requiring selection of contextual cues for particular frames. More recently, Gil, Luciano, Ruiz and Valdivia-Salas (2014) extended their initial work by showing additional patterns of derived hierarchical relations and by providing an improved set of controls over participants' performance. Both these studies have implications for researchers at both the basic and applied level who are potentially interested in understanding and teaching hierarchical responding.

While highly innovative and useful pioneering studies of the examination of hierarchical responding as hierarchical relational framing, both Gil et al. (2012) and Gil et al. (2014) modelled hierarchical responding as broadly conceptualized as opposed to isolating particular subtypes of this behavior. As suggested above, there appear to be a number of varieties of hierarchical responding including, for example, hierarchical classification and hierarchical analysis, and these varieties may have

functionally important differences. Indeed, evidence for this contention comes from mainstream cognitive developmental research. Markman and colleagues (e.g., Markman & Seibert, 1976) compared classification type hierarchy with analysis type hierarchy and found that younger children more readily showed class-inclusion with questions based on the latter than on the former; accordingly, they argued that classification hierarchy develops later than analysis hierarchy. Though cognitive developmental research such as this can provide useful indicators as to the properties of different varieties of hierarchical responding, such work tends to neglect the identification of environmental variables that can facilitate influence over in addition to merely prediction of behavior. This is unsatisfactory for behavioral scientists for whom practical application is critical (see Margolis & Laurence, 2000; Murphy, 2002; Palmer, 2002). In contrast, RFT, as a behavior analytic approach, can suggest a path forward in this regard.

One recent RFT study by Slattery and Stewart (2014) has used relational framing to model hierarchical classification. Mainstream research has shown that hierarchical classification has three core features. *Transitive class containment* refers to classifying a stimulus (A) as a member of a higher order class (C) on the basis that it is a member of a subclass (B) that is a member of that higher order class. For example, if a child is taught that ‘poodle’ (A) is a type of ‘dog’ (B) then they may also classify ‘poodle’ as an ‘animal’ (C) on the basis that ‘dog’ (B) is a type of animal. *Asymmetrical class containment* refers to the fact that a higher order class (e.g., ‘animal’) contains a lower order class (e.g., ‘dog’) but not vice versa (i.e., ‘dog’ does not contain ‘animal’). *Unilateral property induction* refers to the concept that properties or features of a higher order class (e.g., ‘animal’) will also be found in a lower order class (e.g., ‘dog’) but not vice versa. For example, all animals breathe and

thus dogs breathe; however, while dogs have four legs, not all animals do (e.g., Halford, Andrews & Jensen, 2002). The aim of Slattery and Stewart (2014) was to model hierarchical classification as relational framing and examine performance for these properties.

The first part of Phase 1, which established contextual cues, involved conditional discriminative training in which sets of shapes with particular physical features had to be chosen in the presence of particular arbitrary nonsense trigrams (i.e., a three-letter sequence; hereafter coded using alphanumerics; e.g., circles \rightarrow H1; squares \rightarrow H2; blue circle \rightarrow H1.1; yellow circle \rightarrow H1.2; blue square \rightarrow H2.1; yellow square \rightarrow H2.2). This trained participants to relate particular classes of shapes with particular trigram 'labels'. The second part trained them to relate these labels in particular ways in the presence of arbitrary shapes in order to establish the latter as contextual cues for 'member of' and 'includes' respectively. For instance, given 'member of', and with H1.1 as sample, selection of H1 was correct, while given 'includes', and with H1 as sample, selection of H1.1 was correct. Then in Phase 2, these cues were used to train and test a hierarchical relational network of novel trigrams. This involved training particular relations in the network and then testing for mutual and combinatorial entailment and TOF. Results were that ten out of thirteen participants reached the end of Phase 2, with nine exhibiting all three of the properties of hierarchical classification by showing asymmetrical mutual entailment, transitive combinatorial entailment and unidirectional TOF.

The current study aimed to extend this work by using a similar protocol to model hierarchical analysis rather than hierarchical classification and to examine whether the properties of this pattern of relational framing would differ from those implicated in hierarchical classification. The crucial difference between this study and

the previous one was in Phase 1, which established contextual cues. In Slattery and Stewart (2014), the relations that were relevant as regards the training of the contextual cues were between concepts based on the abstraction of common physical properties ('classes') and examples of shapes that included those particular properties ('members'); for instance, between the concept 'green' and particular shapes that were green in colour. In the current study, in contrast, the relations that were relevant as regards the training of the contextual cues were between shapes made up of a number of different parts ('wholes') and examples of the parts themselves ('parts'). For instance, one example of a 'whole' was a compound shape that was comprised of three 'parts' including a green rhombus, a pink arc and a yellow triangle.

As in the previous study, Phase 1 established the functions of contextual cues and then in Phase 2, these cues were used to train and test a hierarchical relational network of previously unseen trigrams, which again involved training particular relations in the network and then testing for mutual and combinatorial entailment and TOF. It was predicted that, because the relational pattern involved was still hierarchical, some properties of the relational framing might be similar to those seen in Slattery and Stewart. However, because this was a different type of hierarchical relational responding (i.e., analytic as opposed to classificative), it was thought possible that other properties might differ. For example, whereas the transfer of function within the Slattery and Stewart model of hierarchical classification had been unidirectional (i.e., from member to class but not from class to member) it was not as clear whether a similar pattern might be seen within hierarchical analysis.

Method

Participants

Ten experimentally naïve participants (six women and four men aged between 21 and 29; $M = 22.6$; $SD = 2.47$) were recruited through personal contacts of the third author. Participants were students at the institution of the corresponding author. None had knowledge of RFT or other forms of stimulus-stimulus relations research. This research was conducted with the formal approval of the institutional ethics committee. All participants signed a statement giving informed consent to the inclusion of material pertaining to themselves, acknowledging that they could withdraw from the study at any time, that they would not be identifiable via the paper; and that they would be fully anonymized.

Apparatus and Materials

Each participant was tested individually in a cubicle containing a desk, a chair and a Fujitsu Siemens™ Scenic PC with a 17" monitor. Instructions, stimulus presentation and recording of responses were controlled by the computer, which was programmed in Visual Basic™ 2008.

Computer-generated stimuli. The stimuli used in Phase 1 (pre-training of non-arbitrary relations) included two arbitrary shapes in black and white as contextual cues (Figure 1, upper panel); 20 stimuli both simple and compound (Figure 1, lower panel); 20 three letter nonsense trigrams (e.g., VEK; henceforth referred to as trigrams); and one blue asterisk. All the shapes used were roughly 1-2 inches square. In Phase 2 (arbitrary relational training and testing) 17 additional trigrams were used.

Insert Figure 1 about Here

Procedure

All participants were trained and tested in sessions lasting between 1.5 and 2 hours. The procedure consisted of: Phase 1, non-arbitrary relational training and testing; Phase 2, arbitrary relational training and testing. Phases comprised various stages of training (with feedback for correct and incorrect responding), and testing (no feedback). All stages were preceded by onscreen instructions to observe the stimuli that would appear and to choose a comparison using the mouse.

Phase 1: Establishing contextual cues. This phase aimed to establish two arbitrary shapes as cues for ‘member of’ and ‘includes’ relational responding, respectively. It included three stages, each with two substages. These three stages, which each used a different set of both shapes and trigrams, used the same basic (conditional discriminative type) format to train the same pattern of contextual control; at the same time, the two later stages also extended their predecessors in particular ways. The rationale behind this procedure, which was developed based on pilot testing and might be considered multiple exemplar training, was to increase the likelihood of appropriate control in Phase 2.

Stage 1. The first substage (1a) involved 48 trials and trained participants to choose particular trigrams in the presence of particular shape stimuli (both simple and compound – see Figure 1). On each trial one stimulus acted as sample and was presented in the top middle center screen. After 1s, two trigrams from the six used in Stage 1 were presented bottom left and bottom right of the screen. After one was chosen, the screen cleared and feedback (‘Correct’ or ‘Incorrect’ in Tahoma, size 44, blue font) was presented on screen for 1s. Then the screen cleared for 1s before the start of the next trial.

The stimuli were non-arbitrarily interrelated along particular physical dimensions such that they could be responded to as ‘wholes’ (compound stimuli) and

‘parts’ (elements within the compounds). More specifically, stimulus set 1 included two simple shapes (i.e., an inverted yellow triangle and a red circle-line shape) and two compound shapes that were made up of juxtapositions of simple shapes (i.e., one compound containing an inverted yellow triangle, a green rhombus and a pink arc and a second compound containing an orange star, a blue rectangle and a red circle-line shape; see Set 1 in Figure 1).

Insert Table 1 About Here

In the first 8 trials of Substage 1a, participants had to choose either of two trigrams, H1 or H2, depending on the sample. H1 had to be chosen in the presence of the triangle-rhombus-arc compound (A1B1C1), while H2 had to be chosen in the presence of the star-rectangle-circle-line compound (A2B2C2). The two trial types (see Table 1, 1a, 1-2) were presented four times each in quasi-random order. In the next 16 trials, participants had to choose one of four additional trigrams (H1.1, H1.2, H2.1, H2.2) in the presence of one particular shape (i.e., triangle (A1) → H1.1; rhombus (B1) → H1.2; star (A2) → H2.1; circle-line (B2) → H2.2). The four trial types (Table 1, 1a, 3-6) were presented four times each in quasi-random order. The final 24 trials presented all six trial types (Table 1, 1a, 1-6) four times each in quasi-random order. The pass criterion was 80% correct across all trials in 1a or 80% correct in the final block of 24. Failure entailed re-exposure up to a maximum of three times.

The aim of Substage 1b was to use the trigrams from 1a to establish arbitrary shapes as cues for hierarchical relations (i.e., ‘part of’ and ‘includes’) by training participants to relate particular trigrams to each other in the presence of those shapes.

On each trial of 1b, a trigram sample was presented in top center screen. After 1s, one of the shapes to be established as contextual cues was presented in middle center screen. Then, 1s later, a number of comparison trigrams (2 or 3, depending on the trial) appeared in a row near the bottom of the screen.

The training of the two contextual cues which was the aim of this substage relied to some extent on the fact that particular interrelated sets of functions had been established in particular trigrams in Substage 1a. For example, H1 had been chosen in the presence of the triangle-rhombus-arc compound; H1.1 had been chosen in the presence of a triangle; and H2.1 had been chosen in the presence of a star. On this basis, the stimulus conditioned with H1 included the stimulus conditioned with H1.1 but not the stimulus conditioned with H2.1. As such, in the presence of H1 as the sample stimulus and the arbitrary shape to be established as the cue ‘includes’, choosing H1.1 rather than H2.1 was correct; while in the presence of H1.1 as sample and the shape to be established as ‘part of’, choosing H1 was correct (not H2).

There were 36 trials presented in a predetermined sequence. The first block of 12 (Table 1, trial types 1b, 1-12 in quasi-random order) established control by the ‘includes’ cue. The second block of 12 (Table 1, trial types 1b, 13-24 in quasi-random order) established control by the ‘part of’ cue. The final block of 12 was a quasi-random mix of both types of trials. Some trials in this and subsequent ‘b’ substages included a blue asterisk, coded ‘X’ in Table 1. On trials in which no other comparison was correct, ‘X’ was deemed correct. Participants had to achieve 92% correct to pass. If they failed they were recycled through 1a before re-exposure to 1b, and this could happen up to a maximum of three times before their participation ended (none of the participants ended in this or subsequent phases because of failure to meet the training criteria).

Stage 2. Stage 2 aimed to test whether the training format in Stage 1 would result in appropriate contextual control with a new stimulus set. This stage was similar to Stage 1 in that it was composed of analogous a and b substages. Substage 2a was identical to 1a except that novel stimuli were used (see Set 2, Figure 1 and 2a, Table 2). If participants passed 2a then they proceeded to 2b. Otherwise they were recycled through 2a again up to a maximum of three times before their participation ended. Substage 2b was similar to 1b but with a number of important differences including the absence of feedback, use of only 24 trials and some differences in trial types (see Table 2, 2b, 1-24). Participants had to achieve 92% correct to progress to Stage 3. Failing this they were recycled back through 2a before re-exposure to 2b. If they again failed 2b then they were re-exposed again but this time with feedback. If they still did not achieve criterion then their participation was ended.

Insert Table 2 About Here

Stage 3. The aim of Stage 3 was to probe for responding in the context of three hierarchical levels rather than just two. Substage 3a was similar to 2a but used novel stimuli (Set 3, Figure 1) and a greater number of stimuli including one compound stimulus that included two other compounds, each of which themselves included two simple shapes, thus constituting a ‘three tier’ part-whole hierarchy. More specifically, stimulus set 3 included four simple shapes (i.e., a yellow quadrilateral, a purple rhombus, a blue pentangle and a pink partial circle), two compound stimuli each made up of two simple shapes (i.e., one containing a yellow quadrangle and a purple rhombus, and a second containing a blue pentagram and a red partial circle) and a third compound stimulus that included the other two compounds

(i.e., the yellow quadrangle and purple rhombus, and the blue pentagram and red partial circle).

Stage 3a involved 64 trials. The set of stimuli used (Set 3, Figure 1) was analogous to the sets used in 1a and 2a in that it included a number of stimuli that were non-arbitrarily interrelated along particular physical dimensions such that they could be responded to as ‘wholes’ (compound stimuli) and ‘parts’ (elements within the compounds). However, because Stage 3 probed for responding with three rather than two hierarchical levels, Set 3 also included one additional comparison stimulus (i.e., a compound containing an orange circle and blue cross). This allowed an additional level of grouping of the shapes in this set, thus enabling additional levels of training in substage 3b. This added complexity also required an expanded set of trigrams (coded H1, H1.1, H1.2, H1.1.1, H1.1.2, H1.2.1, H1.2.2; H2).

Insert Table 3 About Here

In the first 8 trials of 3a, trial types 3a (1-2) (see Table 3, 3a) were quasi-randomly presented four times each. In the next 8, trial types 3a (3-4) were quasi-randomly presented four times each. In the next 16, trial types 3a (5-8) were quasi-randomly presented four times each. In the final 32, a quasi-random mix of all eight previous trial types was presented. Participants had to achieve 80% correct across all trials or 80% correct in the final block of 32 to pass 3a. Failing this they were re-exposed up to three times.

Substage 3b (see Table 3, 3b) was analogous to 2b. There were 24 trials without feedback. Participants had to achieve 92% correct across all trials to pass. Failing this

they were recycled through the entire non-arbitrary training phase before re-exposure to 3b.

Phase 2: Arbitrarily applicable relational training and testing. The aim of this phase was to use the cues for hierarchical relations established during Phase 1 to train an arbitrary hierarchical relational network and to test for derivation of relations and TOF based on that network (see Figure 2). As in Phase 1, a conditional discriminative type format was used throughout.

Insert Figure 2 about Here

Stage 1. Arbitrary relational training. This stage trained arbitrary relations between a number of novel trigrams using the two previously established contextual cue stimuli. On each trial, a trigram sample was presented in top center screen. Following a 1s delay, one of the two cues was presented in middle center screen. Then, 1s later, a number of further trigrams appeared in a row near screen bottom. The participant had to choose one of these. Following selection, the screen cleared and onscreen feedback was presented for 1s. Then the screen cleared for 1s before the next trial.

Six trigrams (H1, H1.1, H1.1.1, H2, H2.1, H2.1.1) were employed as relata. Table 4 (Stage 1) shows the four arbitrary relational training trial types. These were quasi-randomly presented 10 times each in a 40 trial block. Pass criterion was 88% correct in a block or 20 consecutively correct. If a participant passed, they advanced to arbitrary relational testing. If not they were recycled back into arbitrary relational training up to a maximum of three times.

Stage 2. Arbitrary relational testing. This stage probed for the emergence of derived hierarchical relations (see Figure 2). Twelve trial types were used (see Table 4, Stage 2). Trial types 1-6 probed for mutual entailment / asymmetrical class containment, trial-types 7-10 for combinatorial entailment / transitive class containment and trial types 11-12 for a combination of both. These 12 trial-types were presented twice each without feedback in quasi-random order. A minimum of 92% correct (22/24) was needed to progress. Failure to meet criterion meant recycling back through arbitrary relational training and testing up to a maximum of three times.

Insert Table 4 About Here

Stage 3. Training stimulus functions. The aim of this phase was to establish functions for two of the stimuli in the network of trained and derived relations established by the previous stages. In order to investigate whether differences in the position in the relational network in which the functions were established might affect TOF, for half of the participants (1-5), the functions were trained in H1.1 and H2.1 respectively, while for the other half (6-10), they were trained in H1.1.1 and H2 respectively. To accomplish the training, a two choice match to sample (MTS) procedure similar to that used by Gil et al. (2012) was used. On a given trial one of the two stimuli in which functions were being trained acted as sample and was presented in top center screen for 1.5s. Subsequently, two comparison images (one showing grey flecks [F1] and the other blue spikes [F2]) appeared bottom right and left of the screen. Selection of F1 in the presence of H1.1/H1.1.1 and F2 in the presence of H2.1/H2 were deemed correct responses and feedback was similar to previous training phases. There was a total of 16 trials, 4 blocks of 4 trials each with

both samples presented twice per block and with comparison positions counterbalanced. If participants achieved 88% correct they progressed to TOF testing. If not they were recycled back through training up to a maximum of three times.

Stage 4. Testing TOF. This final phase probed for changes in the functions of stimuli in the putative hierarchical relational network established previously, based on the training of, for pts 1-5, F1 in H1.1 and F2 in H2.1 and, for pts 6-10, F1 in H1.1.1 and F2 in H2. As explained in the introduction, it was unclear whether one or more patterns of transformation of functions, if any, might predominate. Possible patterns included (i) downward unidirectional transfer in which there was transfer from stimuli higher up in the hierarchy to stimuli lower down but not from lower to higher (e.g., from H1 to H1.1.1 but not vice versa); (ii) bi-directional transfer both from higher to lower and lower to higher (e.g., from H1 to H1.1.1 as well as vice versa); (iii) upward unidirectional transfer from lower to higher but not higher to lower (e.g., from H1.1.1 to H1 but not vice versa); and (iv) absence of transfer of function in either direction.

TOF testing involved 28 trials in total. Seven stimuli including the six trigrams from previous stages (H1, H1.1, H1.1.1, H2, H2.1, H2.1.1) and one additional (novel) trigram (NT) were used as samples. There were four blocks of seven trials each, and within each block each sample was presented once in a random order. On each trial the sample appeared first in the middle left screen. Next, after 1.5 s, one of two 'feature' stimuli (F1 or F2) was presented in the middle centre screen to the right of the sample. Then, after a further 1.5 s, three response options 'TRUE', 'FALSE', and 'CANNOT SAY' were presented in the top, middle, and bottom right of the screen. After the response the screen cleared for 1.5 s before the next trial.

If participants chose the same option within a trial type for both exposures to that trial type then this was classified as stable responding for that trial type. If they

showed stable responding for all trial types in the test then the results of that exposure were taken as final. If they showed unstable responding for any trial type upon first exposure to the test then they were re-exposed to TOF training and testing once more. If they continued to show instability on this second re-exposure then they were re-exposed to arbitrarily applicable relational responding training and testing with a new stimulus set before being re-exposed to TOF training and testing. They would be allowed a maximum of two exposures to testing in this second session after which the experiment would end even if stable responding had not been shown.

Results

Table 5 shows percentage correct responding for all ten participants. All ten successfully completed both phases.

Insert Table 5 About Here

Phase 1: Establishing contextual cues. Seven out of the ten participants (P1, P2, P6-10) met the training criterion on their first exposure to each of the stages in Phase 1. P3, P4 and P5 failed to meet criterion for 1b on their first attempt (58%, 38% and 52% correct respectively) and thus were recycled through 1a before re-exposure to 1b which they each passed on their second attempt (91%, 98% and 91% correct, respectively). No other re-exposures were needed in Phase 1.

Phase 2: Arbitrarily applicable relational training and testing. All ten participants passed through Stages 1-3 on their first attempt. In TOF testing all ten participants showed consistent responding on first exposure and thus that data was taken as final.

The data for the TOF test are shown in detail in Table 6. Columns 2-9 show the stimuli chosen by each participant for each trial type while Column 10 classifies the pattern of responding. As noted previously, two slightly different forms of transformation of function training and testing were carried out, one for P1-5 and another for P6-10, and thus classification of the pattern of transformation of functions seen depended on which group a participant was in. As such, we will discuss the patterns seen by group.

For P1-5, functions F1 and F2 were established in H1.1 and H2.1, respectively. As such, for this group, in the case of both the trained functions, they could transfer either up or down the hierarchy and thus two patterns could be discerned, one for each function. As suggested previously, possible patterns included (i) downward unidirectional transfer (from stimuli higher up in the hierarchy to stimuli lower down but not from lower to higher – this is the pattern expected for hierarchical classification; for example, imagine that H1 is the superordinate class level ‘Animal’, H1.1 the intermediate class level ‘Dog’ and H1.1.1 the member level ‘Poodle’. If I learn that dogs (H1.1) have ‘grey flecks’ (F1) then I might derive that Poodles (H1.1.1) will also have ‘grey flecks’ (F1) but I would be unable say that all animals (H1) would; this pattern is coded ‘+Uni’ in Table 6); (ii) bi-directional transfer (both from higher to lower and lower to higher; coded ‘Bi’); (iii) upward unidirectional transfer (from lower to higher but not higher to lower; the opposite pattern from hierarchical classification; coded ‘-Uni’); and (iv) absence of transfer of function in either direction (coded ‘None’). As can be seen in Table 6, P1 showed no pattern of TOF; P3 showed upward unidirectional transfer for both functions; P4 and P5 showed bidirectional transfer for both functions; and P2 showed an inconsistent pattern across the two functions involving upward transfer for F1 and bidirectional

transfer for F2. Thus in summary, this group showed little or no evidence of downward transfer and no alternative predominant pattern of transfer.

For P6-10, functions F1 and F2 were established in H1.1.1 and H2 respectively. As such, for this group, the test was whether the trained function would transfer up from H1.1.1 or down from H2 and the combination of these two results would yield one of the same four patterns described in the previous paragraph. Since in this case the pattern could only be determined on the basis of transfer of both functions, only one outcome is reported for each participant in this group. As can be seen in Table 6, P6 and P9 showed bidirectional transfer; P8 showed no pattern of TOF; and P7 and P10 showed unidirectional transfer, with P7 showing downward transfer and P10 upward transfer. Hence, this group showed a similar overall pattern of results as the previous group; there was little or no evidence of downward transfer and no alternative predominant pattern of transfer.

Insert Table 6 About Here

Discussion

The aim of this study was to model and investigate hierarchical analysis or part / whole responding as contextually controlled hierarchical relational responding or relational framing. Ten participants were trained and tested for part / whole hierarchical relations. Similar to participants in Slattery and Stewart (2014) who were trained and tested for hierarchical classification or member / class responding, all ten participants readily showed patterns of framing characterised by asymmetrical mutual entailment and transitive combinatorial entailment, which cognitive developmental theorists have argued are important characteristics of hierarchical responding. With

respect to the transformation of functions, however, participants in the current study showed a substantially different pattern of responding than that seen in Slattery and Stewart (2014). Whereas in the latter, almost all participants showed downward unidirectional TOF, in the current study, only one out of ten showed this pattern, while the remaining nine showed a variety of other patterns including bidirectional TOF in the case of four; upward unidirectional TOF (i.e., the opposite pattern to downward unidirectional TOF) in the case of two; no TOF in the case of two; and an inconsistent pattern in the case of the remaining participant.

As regards both the patterns of mutual and combinatorial entailment, which were the same as in Slattery and Stewart (2014), and the pattern of TOF, which showed a substantial difference from that seen in that study, the results found are in accordance with prediction to an important extent. It had been predicted that this study would demonstrate patterns of asymmetrical mutual and transitive combinatorial entailment just as in Slattery and Stewart because these are features of all hierarchical responding. Meantime, it had been suggested that the pattern of TOF might differ from that seen in the previous study, because the current study was investigating what previous evidence had suggested was a functionally different type of hierarchical responding.

As outlined in the introduction, cognitive developmental researchers have previously provided evidence of hierarchical analysis or part / whole responding as being a functionally different pattern of responding from hierarchical classification or member / class responding. For instance, Markman and colleagues (e.g., Markman & Seibert, 1976) compared member / class ('class-concept') hierarchy with part / whole ('collection-concept') hierarchy and found that younger children more readily showed class-inclusion with part / whole than with member / class questions, suggesting that

the latter develops later than the former. However, while research such as this constitutes empirical evidence for two different patterns of hierarchical responding, from a behavior analytic perspective, we still require a bottom up functional analytic account of the processes involved in these two different patterns.

Slattery and Stewart (2014) and the present study are part of a research stream whose aim is to investigate these different forms of hierarchy using such a bottom up behavior analytic approach. More specifically, these studies have used an RFT approach by modelling hierarchical responding as contextually controlled hierarchical relational responding or relational framing. As explained earlier, in RFT, relational framing is seen as the key psychological process involved in language and cognition. Relational framing is arbitrarily applicable relational responding, in which the relational responding is primarily under contextual control that determines the relation as opposed to being under the control of physical or non-arbitrary properties of the related stimuli. For example, in hierarchical relational framing, cues such as ‘part of’ or ‘includes’ come to control a pattern of relational responding that can be applied even in the absence of actual physical properties. For instance, if I am told that ‘object X is part of object Y’, then, without seeing either object or being told anything else, I might derive that ‘object Y includes object X’. However, according to RFT, patterns of abstract relational framing such as this have their roots in non-arbitrary relational responding, in which stimuli are related based on physical properties. For example, children initially learn to discriminate things as being physically part of other things before they come to be able to respond in accordance with such relations at a more abstract level. The concept underlying the current research stream is that, since non-arbitrary relational roots can determine the pattern of abstract relational responding, then one important source of the difference between part-whole and member-class

hierarchical relational framing might be found at the non-arbitrary level and thus one way in which to model and compare these two forms of hierarchical relational responding might be by using particular non-arbitrary relational training protocols.

This is exactly what has been done in Slattery and Stewart (2014) and the current study. In the former, the aim of the non-arbitrary relational training protocol was to train participants to respond to stimuli as part of a collective based on shared features (i.e., a class of which they are members). In the current study, in contrast, the non-arbitrary relational training aimed to train participants to respond to stimuli as part of a collective based instead on proximity (i.e., a whole of which they are part). It was predicted that whereas the patterns of mutual and combinatorial entailment for the two protocols would be similar, the patterns of TOF might differ. The latter was predicted partly because of the relationships seen at a non-arbitrary level. In the non-arbitrary training for member / class hierarchical responding, the collective is based on shared features and so any feature of the class should (at least logically) be shared by members. At the same time, not every feature of the individual members of the class will be a feature of the collective because otherwise all class members would be identical. Hence, in new contexts that feature the cues trained on the basis of this type of training, TOF from top (class) to bottom (member) might be relatively likely and TOF from bottom (member) to top (class) relatively unlikely. Meantime, in the non-arbitrary training for part / whole hierarchical responding, the collective is based on proximity (grouping) of the parts and thus these same relationships do not apply and thus in a context involving contextual cues trained up on this basis it is less clear what functions might transfer.

One possible critique of the current study might focus on the assumption of proximity as a defining element of the part-whole relation. It might be argued that

perhaps other features might also support a part-whole relationship. For example, in the case of the ‘nail/finger/hand’ example it may not just be proximity which supports the abstraction of a part whole relationship – it could also be colour, shape etc. A related point that might be made in this regard is that perhaps altering the characteristics of some of the exemplars might yield less variability in the transformation of functions. For example, perhaps it could be argued that variation of physical characteristics of the parts so as to be more similar to one another might have yielded a more consistent pattern of transformation of functions and perhaps one more similar to that seen in Slattery and Stewart (2014). However, while it may indeed be true that features such as physical similarity might support the abstraction of a part-whole pattern, we would suggest that proximity is still the central and defining aspect of the part-whole relation. As such, for the current study we aimed to isolate that aspect in particular when examining for the possibility of a different pattern of transformation of functions than seen in Slattery and Stewart (2014).

Assuming that proximity is indeed a key aspect of the non-arbitrary training that underlies part-whole relations, this non-arbitrary relational influence is likely compounded to some extent by historical experience of the socio-verbal contingencies associated with the two sets of contextual cues (part / whole and member / class) at issue. In cases involving member / class hierarchical classification, the pattern of TOF from class to member but not vice versa is often specified as conforming with a downward unidirectional transfer. Contexts in which such specification might occur might include episodes of scientific education or discussion, to which at least some and possibly all of the Western educated adults in both Slattery and Stewart and the current study have likely been exposed at various points in their school or university training. Meantime, as regards hierarchical analysis or part-whole relations, however,

analogous types of experience that might have trained similarly well specified relations seem less likely and common experience with such relations would likely not provide such specification either.

Consider an example. If I'm part of a prestigious organization then that may raise my reputation (transfer of status from whole to part); similarly, if I already have a positive reputation before it is known that I'm a part of a particular group then knowledge of my status might boost the group's reputation (part to whole). Perhaps in the former case there might be more transfer than in the latter case so that though there is bidirectional transfer it is asymmetrical in "strength". There are also likely many exemplars of good people working for bad organisations and vice versa and so based on those examples there might be less transfer (or even an absence of transfer) in either direction. This suggests that in the less well defined arena of part-whole relations, aspects of context might matter more in determining the pattern of TOF. In any event, it seems less clear what form of TOF might be seen and this fact alone might lead one to predict that at the very least a universal downward unidirectional transfer of function is relatively improbable and thus this pattern would differ from that seen in the case of hierarchical classification.

The results seen in the current experiment support the hypothesis that the TOF for classification and part-whole hierarchy might differ. Furthermore, they support the idea of a less well specified pattern for the latter than the former because, apart from not seeing a predominantly downward transfer of function for part-whole hierarchy, there was in fact no pre-dominant pattern of any kind seen. Instead, there was a mixture of different patterns with none in the majority. This is a useful first step in the investigation of part-whole hierarchy. Future research looking at this phenomenon might attempt to examine it in further depth. For example, might one or other patterns

be made more likely to predominate in certain conditions? For instance, perhaps participants might be primed with particular examples of part-whole relations before exposure to the model in order to examine the extent to which transfer might be influenced by aspects of context.

One previous study that is relevant with regard to manipulation of contextual control is Wulfert and Hayes (1988) who demonstrated the manipulation of contextual control over transfer of function through equivalence relations. In this study, participants first showed transfer of an ordering response through derived equivalence relations and then subsequently, the direction of the ordering was brought under contextual control (specifically, under the control of different tones). In the case of Wulfert and Hayes, the relations involved were equivalence or coordinate relations and thus this was a simpler pattern of relational responding in which (within an RFT approach) a transfer of some kind would be expected to be well established and specified based on an individual's history of exposure to the verbal community. Nevertheless, this work is suggestive of the type of experimental manipulation that might be brought to bear in a protocol such as the present one, for example, during or after training and testing for derived hierarchical relations, so as to specify stimulus control over particular patterns of subsequent transformation of function or indeed, for the absence of any obvious pattern. This is work in which future research into hierarchical relational framing might engage.

A further possible future direction for research might be to use the current model to explore the development of hierarchical responding. As described, previous cognitive developmental research has examined hierarchical responding in children and has found that part-whole responding appears to emerge sooner than member / class responding. This work was one source of impetus for the current research. As

has been suggested, however, a key advantage of behavior analytic work such as the current research is that it emphasizes the identification of variables that might facilitate influence over behavior. The current research suggests that an RFT approach to hierarchical responding in terms of contextually controlled relational responding might prove fruitful. For example, initial cross sectional research might explore at what ages children's behavior tends to come under the control of various cues for hierarchical responding. This work might be used to develop a variety of forms of tasks varying in complexity, abstractness and type of hierarchical relation. Subsequently, protocols refined on the basis of such work might be used to assess and train children showing deficits with respect to this repertoire. The latter work might be conducted first with typically developing children whose deficits are more likely related to age and subsequently with children with educational or developmental delays.

In conclusion, hierarchical responding is a key repertoire important both in basic decision making as well as advanced problem solving and analysis (e.g., critical thinking, scientific analysis). As such, the development of effective procedures for exploring, assessing and training these repertoires in both typically developing and developmentally delayed populations are important goals. The current work is part of a broader stream of RFT research investigating hierarchical responding. It is hoped that this work is taking useful steps in the direction of greater theoretical and practical understanding of the repertoires involved.

Statement of Ethics Issues

The studies reported herein have been approved by the National University of Ireland Research Ethics Committee. They have been performed in accordance with the ethical standards as laid down in the 1964 Declaration of Helsinki and its later amendments.

Conflict of Interest Statement

The authors report no conflicts of interest.

References

- Dougher, M.J., Hamilton, D.A., Fink, B. C., & Harrington, J. (2007). Transformation of the discriminative and eliciting functions of generalized relational stimuli. *Journal of the Experimental Analysis of Behavior*, *88*(2), 179-197.
DOI: 10.1901/jeab.2007.45-05.
- Dymond, S., & Roche, B. (2013). *Advances in Relational Frame Theory: Research and Application*. CA: New Harbinger Publications, Inc.
- Gil, E., Luciano, C., Ruiz, F. J., & Valdivia-Salas, S. (2012). A preliminary demonstration of transformation of functions through hierarchical relations. *International Journal of Psychology and Psychological Therapy*, *12*(1), 1-19.
- Gil, E., Luciano, C., Ruiz, F. J., & Valdivia-Salas, S. (2014). Towards a functional analysis of hierarchical classification: A further experimental step. *International Journal of Psychology and Psychological Therapy*, *14*(2), 147-153.
- Griffee, K., & Dougher, M. J. (2002). Contextual control of stimulus generalisation and stimulus equivalence in hierarchical categorization. *Journal of Experimental Analysis of Behavior*, *78*(3), 433-447. DOI: 10.1901/jeab.2002.78-433.
- Halford, G.S. Andrews, G., & Jensen, I. (2002). Integration of category induction and hierarchical classification: One paradigm at two levels of complexity. *Journal of Cognition and Development*, *3*(2), 143-177. DOI: 10.1207/S15327647JCD0302_2.
- Hayes, S. C., Barnes-Holmes, D., & Roche, B. (2001). *Relational Frame Theory: A post-Skinnerian account of human language and cognition*. NY: Kluwer Academic / Plenum Publishers.

- Margolis, E., & Laurence, S. (2000). *Concepts: Core Readings*. Cambridge, MA: The MIT Press.
- Markman, E. M., & Seibert, J. (1976). Classes and collections: Internal organization and resulting holistic properties. *Cognitive Psychology*, 8, 516-577.
- Murphy, G. (2002). *The Big Book of Concepts*. Cambridge, MA: The MIT Press.
- Palmer, D. C. (2002). Psychological Essentialism: A review of E. Margolis and S. Laurence (Eds.), *Concepts: Core Readings*. *Journal of the Experimental Analysis of Behavior*, 78(3), 597-607. DOI: 10.1901/jeab.2002.78-597.
- Slattery, B., & Stewart, I. (2014). Hierarchical classification as relational framing. *Journal of the Experimental Analysis of Behavior*, 101(1), 61-75. DOI: 10.1002/jeab.63.
- Slattery, B., Stewart, I., & O'Hora, D. (2011). Testing for transitive class containment as a feature of hierarchical classification. *Journal of the Experimental Analysis of Behavior*, 96(2), 243-260. DOI: 10.1901/jeab.2011.96-243.

Table 1. Trial types in Phase 1 (establishing contextual cues) Stage 1. ‘I’ and ‘P’ are the cues ‘Includes’ and ‘Part of’, respectively. The comparisons are separated by commas. X is a blue asterisk which was designated correct when no other comparison was. Note that all trial types in Phase 1 Stage 1 were training trials (i.e., had feedback).

Ph.	Stg.	No.	Contextual Cue	Sample	Comparisons	Correct Comparison	Feedback
1	1a	1	[None]	A1B1C1	H1, H2	H1	Yes
		2	[None]	A2B2C2	H2, H1	H2	Yes
		3	[None]	A1	H1.1, H1.2	H1.1	Yes
		4	[None]	B1	H1.2, H1.1	H1.2	Yes
		5	[None]	A2	H2.1, H2.2	H2.1	Yes
		6	[None]	B2	H2.2, H2.1	H2.2	Yes
1	1b	1	I	H1	H1.1, H2.1, X	H1.1	Yes
		2	I	H1	H1.1, H2.2, X	H1.1	Yes
		3	I	H1	H1.2, H2.1, X	H1.2	Yes
		4	I	H1	H1.2, H2.2, X	H1.2	Yes
		5	I	H2	H2.1, H1.1, X	H2.1	Yes
		6	I	H2	H2.1, H1.2, X	H2.1	Yes
		7	I	H2	H2.2, H1.1, X	H2.2	Yes
		8	I	H2	H2.2, H1.2, X	H2.2	Yes
		9	I	H1.1	X, H1, H2	X	Yes
		10	I	H1.2	X, H1, H2	X	Yes
		11	I	H2.1	X, H1, H2	X	Yes
		12	I	H2.2	X, H1, H2	X	Yes
		13	P	H1.1	H1, H2, X	H1	Yes
		14	P	H1.1	H1, H1.1, X	H1	Yes
		15	P	H1.2	H1, H2, X	H1	Yes
		16	P	H1.2	H1, H1.2, X	H1	Yes
		17	P	H2.1	H2, H1, X	H2	Yes
		18	P	H2.1	H2, H2.1, X	H2	Yes
		19	P	H2.2	H2, H1, X	H2	Yes
		20	P	H2.2	H2, H2.2, X	H2	Yes
		21	P	H1	X, H1, H1.1	X	Yes
		22	P	H1	X, H1, H1.2	X	Yes
		23	P	H2	X, H1, H2.1	X	Yes
		24	P	H2	X, H1, H2.2	X	Yes

Table 2. Trial types in Phase 1 (establishing contextual cues) Stage 2. ‘I’ and ‘P’ are the cues ‘Includes’ and ‘Part of’, respectively. The comparisons are separated by commas. X is a blue asterisk which was designated correct when no other comparison was. Note that all trial types in Phase 1 Stage 2a were training trials (i.e., had feedback) while those in Phase 1 Stage 2b were testing trials (i.e., had no feedback).

Ph.	Stg.	No.	Contextual Cue	Sample	Comparisons	Correct Comparison	Feedback
1	2a	1	[None]	A3B3C3	H1, H2	H1	Yes
		2	[None]	A4B4C4	H2, H1	H2	Yes
		3	[None]	A3	H1.1, H1.2	H1.1	Yes
		4	[None]	B3	H1.2, H1.1	H1.2	Yes
		5	[None]	A4	H2.1, H2.2	H2.1	Yes
		6	[None]	B4	H2.2, H2.1	H2.2	Yes
1	2b	1	I	H1	H1.1, H2.1, X	H1.1	No
		2	I	H1	H1.1, H1, X	H1.1	No
		3	I	H1	H1.2, H2.2, X	H1.2	No
		4	I	H1	H1.2, H1, X	H1.2	No
		5	I	H2	H2.1, H1.1, X	H2.1	No
		6	I	H2	H2.1, H2, X	H2.1	No
		7	I	H2	H2.2, H1.2, X	H2.2	No
		8	I	H2	H2.2, H2, X	H2.2	No
		9	I	H1.1	X, H1, H1.1	X	No
		10	I	H1.2	X, H1, H1.2	X	No
		11	I	H2.1	X, H1, H2.1	X	No
		12	I	H2.2	X, H1, H2.2	X	No
		13	P	H1.1	H1, H2, X	H1	No
		14	P	H1.1	H1, H1.1, X	H1	No
		15	P	H1.2	H1, H2, X	H1	No
		16	P	H1.2	H1, H1.2, X	H1	No
		17	P	H2.1	H2, H1, X	H2	No
		18	P	H2.1	H2, H2.1, X	H2	No
		19	P	H2.2	H2, H1, X	H2	No
		20	P	H2.2	H2, H2.2, X	H2	No
		21	P	H1	X, H1.1, H2.1	X	No
		22	P	H1	X, H1.1, H2.1	X	No
		23	P	H2	X, H1.1, H2.1	X	No
		24	P	H2	X, H1.1, H2.1	X	No

Table 3. Trial types in Phase 1 (establishing contextual cues) Stage 3(a, b). ‘I’ and ‘P’ are the cues ‘Includes’ and ‘Part of’, respectively. The comparisons are separated by commas. X is a blue asterisk which was designated correct when no other comparison was. Note that all trial types in Phase 1 Stage 3a were training trials (i.e., had feedback) while all trial types in Phase 1 Stage 3b were testing trials (i.e., had no feedback).

Ph.	Stg.	No.	Contextual Cue	Sample	Comparisons	Correct Comparison	Feedback
1	3a	1	[None]	A5B5C5D5	H1, H2	H1	Yes
		2	[None]	E1	H2, H1	H2	Yes
		3	[None]	A5B5	H1.1, H1.2	H1.1	Yes
		4	[None]	C5D5	H1.2, H1.1	H1.2	Yes
		5	[None]	A5	H1.1.1, H1.1.2	H1.1.1	Yes
		6	[None]	B5	H1.1.2, H1.1.1	H1.1.2	Yes
		7	[None]	C5	H1.2.1, H1.2.2	H1.2.1	Yes
		8	[None]	D5	H1.2.2, H1.2.1	H1.2.2	Yes
1	3b	1	I	H1.1.1	X, H1, H1.1.1	X	No
		2	I	H1.1.2	X, H1, H1.1	X	No
		3	I	H1.2.1	X, H1, H1.2	X	No
		4	I	H1.2.2	X, H1, H2	X	No
		5	I	H1	H1.1.1, H1, X	H1.1.1	No
		6	I	H1	H1.1.2, H2, X	H1.1.2	No
		7	I	H1	H1.2.1, H1, X	H1.2.1	No
		8	I	H1	H1.2.2, H1, X	H1.2.2	No
		9	I	H2	X, H1.1, H1.2.1	X	No
		10	I	H1.1	H1.1.1, H1, X	H1.1.1	No
		11	I	H1.2	H1.2.1, H1.2, X	H1.2.1	No
		12	I	H1.2	X, H1, H2	X	No
		13	P	H1.1.1	H1, H2, X	H1	No
		14	P	H1.1.2	H1, H2, X	H1	No
		15	P	H1.2.1	H1, H2, X	H1	No
		16	P	H1.2.2	H1, H2, X	H1	No
		17	P	H1	X, H1.1.1, H1	X	No
		18	P	H1	X, H1.2.1, H2	X	No
		19	P	H1	X, H1.1.2, H1.1	X	No
		20	P	H1	X, H1.2.2, H1.2	X	No
		21	P	H1.1	X, H1.1, H1.1.1	X	No
		22	P	H1.1	X, H1.2, H1.1.2	X	No
		23	P	H1.2	X, H2, H1.2.1	X	No
		24	P	H2	X, H2, H1.2.2	X	No

Table 4. Trial types in Phase 2 Stage 1 (arbitrarily applicable relational training) and Phase 2 Stage 2 (arbitrarily applicable relational testing). ‘I’ and ‘P’ are the cues ‘Includes’ and ‘Part of’, respectively. The comparisons are separated by commas. Note that all trial types in Phase 2 Stage 1 were training trials (i.e., had feedback) while all trial types in Phase 2 Stage 2 were testing trials (i.e., had no feedback). N1 and N2 were novel trigrams.

Ph.	Stg.	No.	Contextual Cue	Sample	Comparisons	Correct Comparison	Feedback
2	1	1	I	H1	H1.1, H2.1	H1.1	Yes
		2	I	H1.1	H1.1.1, H2.1.1	H1.1.1	Yes
		3	P	H2.1.1	H2.1	H2.1	Yes
		4	P	H2.1	H2, H1	H2	Yes
2	2	1	P	H1.1	H1, N1, H1.1.1	H1	No
		2	P	H1.1.1	H1.1, H2, H2.1	H1.1	No
		3	I	H1.1	N1, H1, H2	N1	No
		4	I	H2	H2.1, H1.1, H1.1.1	H2.1	No
		5	I	H2.1	H2.1.1, N2, H2	H2.1.1	No
		6	P	H2.1	N1, H2.1.1, H1.1.1	N1	No
		7	I	H1	H1.1.1, N1, H2.1.1	H1.1.1	No
		8	P	H1.1.1	H1, N2, H2	H1	No
		9	P	H2.1.1	H2, N1, H1	H2	No
		10	I	H2	H2.1.1, H2, H1.1	H2.1.1	No
		11	P	H1	N1, H1, H1.1.1	N1	No
		12	I	H2.1.1	N1, H2, H2.1.1	N1	No

Table 5. Percentage correct responding in stages 1-4 of Phase 1 and stages 1-3. Failure to meet pass criteria are highlighted in italics and bold font. An asterisk denotes a stage of training, on which feedback was provided. On testing trials (no asterisk), no feedback was provided on any trial.

Pt	Stg 1		Phase 1				Phase 2		
	1a*	1b*	Stg 2		Stg 3		Stg 1	Stg 2	Stg 3
			2a*	2b	3a*	3b	AAR*	AAR	TOF*
1	98	93	100	96	97	92	98	100	100
2	98	93	98	100	95	92	95	100	100
3	94	96	96	96	95	92	98	92	94
4	83	93	96	100	95	92	88	92	100
5	96	96	96	100	97	96	95	96	100
6	100	91	96	92	88	96	90	96	94
7	90	93	100	100	97	92	98	100	94
8	96	58							
	100	91	98	100	96	92	100	100	94
9	100	38							
	100	98	88	100	96	92	100	96	100
10	96	52							
	100	91	90	96	95	96	100	92	88

Part-Whole Hierarchical Relations

Table 6. Data for transformation of functions testing (Phase 2, Stage 4). Fn. = the trained and tested function flecks’) or F2 (‘blue spikes’); H1, H1.1, H1.1.1, H2, H2.1 and H2.1.1 were nonsense trigrams predicted to be network; NT = Novel nonsense trigram; T = True; F = False; C = Cannot say; +Uni = pattern of responses in unidirectional transformation of functions; -Uni = pattern of responses in accordance with an upward unidirectional transformation of functions; Bi = pattern of responses in accordance with a bidirectional transformation of functions.

Pt.	Fn.	H1	H1.1	H1.1.1	NT	H2	H2.1	H2.1.1
P1	F1	F, F	T, T	F, F	C, C	F, F	F, F	F, F
	F2	F, F	F, F	F, F	C, C	F, F	T, T	F, F
P2	F1	T, T	T, T	F, F	C, C	F, F	F, F	F, F
	F2	F, F	F, F	F, F	C, C	T, T	T, T	T, T
P3	F1	T, T	T, T	C, C	F, F	F, F	F, F	F, F
	F2	F, F	F, F	F, F	F, F	T, T	T, T	C, C
P4	F1	T, T	T, T	T, T	F, F	F, F	F, F	F, F
	F2	F, F	F, F	F, F	F, F	T, T	T, T	T, T
P5	F1	T, T	T, T	T, T	C, C	F, F	F, F	F, F
	F2	F, F	F, F	F, F	C, C	T, T	T, T	T, T
P6	F1	T, T	T, T	T, T	C, C	F, F	F, F	F, F
	F2	F, F	F, F	F, F	C, C	T, T	T, T	T, T
P7	F1	C, C	C, C	T, T	C, C	F, F	F, F	F, F
	F2	F, F	F, F	F, F	C, C	T, T	T, T	T, T
P8	F1	F, F	F, F	T, T	F, F	F, F	F, F	F, F
	F2	F, F	F, F	F, F	F, F	T, T	F, F	F, F
P9	F1	T, T	T, T	T, T	C, C	F, F	F, F	F, F
	F2	F, F	F, F	F, F	C, C	T, T	T, T	T, T
P10	F1	T, T	T, T	T, T	C, C	F, F	F, F	F, F
	F2	F, F	F, F	F, F	C, C	T, T	C, C	C, C

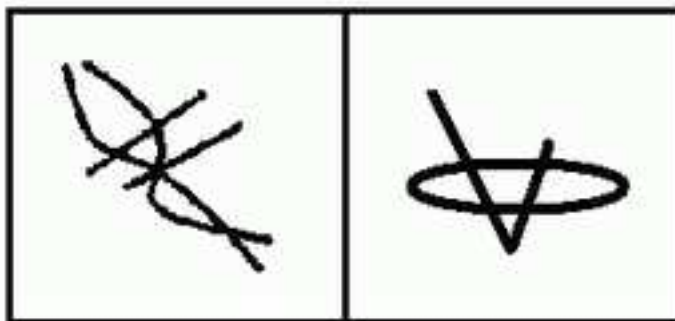
Figure Captions

Figure 1. Upper: The arbitrary stimuli established as contextual cues for ‘part of’ and ‘includes’ relations. Lower: The simple and compound shape stimuli used in the non-arbitrary relational training phases. Set 1 included four simple shapes (A1, B1, A2 and B2) and two compound shapes (A1B1C1 and A2B2C2); Set 2 included four simple shapes (A3, B3, A4, B4) and two compound shapes (A3B3C3 and A4B4C4); and Set 3 included four simple shapes (A5, B5, A6, B6) and four compound shapes (A5B5, C5D5, A5B5C5D5 and E1). The alphanumeric codes H1, H2, H1.1, H1.2, H2.1 and H2.2 represent nonsense trigrams which had to be selected in the presence of particular shapes. Their make-up is intended to ease communication regarding the potential hierarchical relationships that training aimed to engender. Different nonsense trigrams were used in the case of each of the three stages.

Figure 2. Networks of trained and tested relations in Stage 2. Solid arrows indicate trained relations, dashed arrows indicate derived relations. Alphanumerics (i.e., H1, H1.1, H1.1.1; H2, H2.1, H2.1.1) represent the nonsense syllable stimuli used in training and testing. The letters ‘P’ and ‘I’ denote contextually controlled hierarchical relations ‘part of’ and ‘includes’ respectively.

Figure 1

CONTEXTUAL CUES



EXPERIMENTAL STIMULI

SET 1	<i>A1B1C1</i> 		<i>A2B2C2</i> 	
	<i>H1</i>		<i>H2</i>	
	<i>A1</i> 	<i>B1</i> 	<i>A2</i> 	<i>B2</i>
	<i>H1.1</i>	<i>H1.2</i>	<i>H2.1</i>	<i>H2.2</i>
SET 2	<i>A3B3C3</i> 		<i>A4B4C4</i> 	
	<i>H1</i>		<i>H2</i>	
	<i>A3</i> 	<i>B3</i> 	<i>A4</i> 	<i>B4</i>
	<i>H1.1</i>	<i>H1.2</i>	<i>H2.1</i>	<i>H2.2</i>
SET 3	<i>A5B5C5D5</i> 			<i>E1</i>
	<i>H1</i>			<i>H2</i>
	<i>A5B5</i> 		<i>C5D5</i> 	
	<i>H1.1</i>		<i>H1.2</i>	
	<i>A5</i> 	<i>B5</i> 	<i>C5</i> 	<i>D5</i>
	<i>H1.1.1</i>	<i>H1.1.2</i>	<i>H1.2.1</i>	<i>H1.2.2</i>

Figure 2

