UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

From Pirate Islands to Routing Tables: Investigating Intermediate Representations in Concreteness Fading through AR Learning

Trory, A.; Howland, K.; Good, J.; du Boulay, B.

DOI
10.1145/3628516.3655817

Publication date
2024

Document Version
Final published version

Published in
Proceedings of ACM Interaction Design and Children Conference (IDC)

License
CC BY

Link to publication

Citation for published version (APA):
Trory, A., Howland, K., Good, J., & du Boulay, B. (2024). From Pirate Islands to Routing Tables: Investigating Intermediate Representations in Concreteness Fading through AR Learning. In *Proceedings of ACM Interaction Design and Children Conference (IDC): 17-20 June 2024 : Delft NL : "Inclusive happiness"* (pp. 469–479). Association for Computing Machinery. https://doi.org/10.1145/3628516.3655817

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:13 Nov 2026

From Pirate Islands to Routing Tables: Investigating Intermediate Representations in Concreteness Fading through AR Learning

Anthony Trory
University of Sussex
a.trory@sussex.ac.uk

Kate Howland
University of Sussex
k.l.howland@sussex.ac.uk

Judith Good
University of Amsterdam
j.a.good@uva.nl

Benedict du Boulay
University of Sussex
b.du-boulay@sussex.ac.uk

ABSTRACT

With the recent growth in interest in computing education for younger children, concrete representations have enabled children to understand concepts that, in abstract form, may otherwise be considered too advanced. Introducing children to computing concepts through concrete representations may be necessary at younger ages, but there is also a need for effective methods for progressing from this towards abstract understanding. Concreteness Fading has previously been applied in computing contexts with some success, but questions remain about how to design learning environments that most effectively support this progression, including the role of intermediate representations. This paper presents an augmented reality (AR) environment that uses the Concreteness Fading approach to teach introductory topics in computer networks. An in-school experiment with 59 children aged 9-10 investigated the effect of three distinct methods of fading representations from concrete to abstract, using a between-groups pre-test / post-test design. Three groups were compared: concrete > concrete > abstract ("two-step"), concrete > intermediate > abstract ("three-step"), and concrete > concrete/intermediate > intermediate > intermediate/abstract > abstract ("five-step"). The results showed a statistically significant difference between the two- and three-step groups in favour of the three-step, confirming the value of explicitly linking mutual referents via an intermediate representation when moving from concrete to abstract. No statistically significant difference was found between the three- and five-step groups, suggesting that increasing the number of intermediate representations does not significantly improve the learning outcomes. We discuss the implications of these findings for the design of educational technology for children.

CCS CONCEPTS

• **Human-centered computing** → Interface design prototyping; • **Applied computing** → Interactive learning environments; • **Social and professional topics** → K-12 education.

This work is licensed under a Creative Commons Attribution International 4.0 License.

IDC '24, June 17–20, 2024, Delft, Netherlands
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0442-0/24/06
https://doi.org/10.1145/3628516.3655817

KEYWORDS

Concreteness Fading, Computing Education, Educational Technology

ACM Reference Format:
Anthony Trory, Kate Howland, Judith Good, and Benedict du Boulay. 2024. From Pirate Islands to Routing Tables: Investigating Intermediate Representations in Concreteness Fading through AR Learning. In *Interaction Design and Children (IDC '24), June 17–20, 2024, Delft, Netherlands*. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3628516.3655817

1 INTRODUCTION

Concrete representations, such as visualisations, metaphors, and physical manipulatives, have a long history of use in the computing classroom [8.]. However, concrete materials and activities alone are not always effective in computing education [3.], and there is a need for planned progression to more abstract materials [28.]. With roots in Bruner's Modes of Representation [6.], Concreteness Fading describes an approach in which a concept is first introduced using a concrete representation (such as using physical counters when teaching addition) followed by more abstract representations (ending, for example, with a numerical sum). Varying modes of Concreteness Fading have been used in STEM education, most frequently in mathematics, but with numerous applications in science and computing education [25.]. The number and form of steps in Concreteness Fading is not standardised, with many following Bruner's suggestion of three steps, but some adopting two (e.g. [15.]) or four (e.g. [23.]) steps in the fading process. There is growing evidence for the effectiveness of the approach (e.g. [12., 28.]), but contradictory findings are reported [16.]. Given the differing interpretations and implementations, it is perhaps unsurprising that systematic reviews report mixed findings on the efficacy of Concreteness Fading [13., 25.]. There remain open questions, including the importance of physicality in the concrete step, whether it is necessary for intermediate steps to employ explicit linking of mutual referents, and whether it is necessary to have more than two steps at all [14.].

The Theoretical Model of Concreteness Fading (TMCF) provides the most detailed framework for scaffolding the move from concrete to abstract [13.] and has shown promising effects in mathematics [12.] and computing [28.]. The TMCF advocates the gradual perceptual fade from an initial concrete representation towards the abstract over at least three steps, arguing that perceptual richness and three-dimensionality are important at the concrete step and that explicit referencing between each step is more effective than

presenting consecutive representations that do not include mutual referents [14.]. Whilst Fyfe argues that a physical concrete representation may offer the highest level of effectiveness, there is evidence that virtual concreteness can also be effective in Concreteness Fading (e.g. [5., 15.]). Augmented Reality (AR) technology has the potential to support a three-dimensional concrete representation of high perceptual richness, which can gradually fade in-place such that mutual referents are explicitly linked through contiguity.

This paper explores whether Concreteness Fading can be effectively implemented in an AR learning environment for teaching computer network concepts and investigates whether a three-step fading process is more effective than a two-step fading process in this context. In addition, since the TMCF argues that smooth fading is more effective than abrupt shifts [13.], we investigate whether a five-step process of more gradual fading is more effective than a three-step process. We present an AR learning environment designed to support gradual fading and an experiment to evaluate the environment conducted in an authentic school setting. The experiment investigated two research questions and two hypotheses. **Research Question 1**: How does linking the concrete and abstract representations through a single intermediate representation affect learning gains? **Hypothesis 1**: The 3-step group will have higher learning gains than the 2-step group. **Research Question 2**: How does a more gradual fade through an increased number of intermediate representations affect learning gains? **Hypothesis 2**: The 5-step group will have higher learning gains than the 3-step group.

Our contributions are empirical and technical. We provide empirical evidence that, within an AR learning environment, a three-step implementation of Concreteness Fading that includes an intermediate representation in which mutual referents are explicitly linked is more effective than a two-step implementation, while overall time on task was kept constant. We did not find evidence that a more gradual fade over five steps is more effective than fading over three steps. We also make a technical contribution with an AR learning environment that supports in-place fading from a perceptually rich 3D concrete representation to a 2D abstract symbolic representation.

2 RELATED WORK

Concreteness Fading can be understood as a specific approach to designing and integrating multiple external representations (MER) for learning. Ainsworth [1.] argues that studies of MERs show mixed results. However, one consistent finding is that learners often face difficulties translating between representations, suggesting that identifying and linking their mutual referents is not automatic. The author notes that learning environments sometimes attempt to lower the cognitive load required for translation, for example, through "dyna-linking": interacting with representation A updates the display of representation B, etc. Although these findings are primarily from learning environments that display MERs simultaneously, it may be safe to assume the sequential presentation inherent to Concreteness Fading creates a similar or greater burden on learners. Therefore, it is important to increase our understanding of different approaches to linking mutual referents within Concreteness Fading, as is the aim of this study.

Suh et al.'s [25.] systematic literature review on using Concreteness Fading in design for children observes that the existing literature is fragmented and that conflicting results are reported across studies. The authors highlight the variation amongst approaches, including the order of progression, form and quantity of representations, presentation timing and connection between representations. This variation and lack of specificity likely accounts for the mixed results. For example, the authors refer to a study by Butler et al. [7.], comparing two- and three-step approaches to Concreteness Fading. In the two-step group, the sequence began with an "intermediate" representation instead of the usual concrete. The study found a significant increase in learning outcomes for the three-step group, but these results could perhaps be explained by the lack of initial concreteness rather than the number of representations.

One prominent example of two-step Concreteness Fading was the work of Goldstone and Son [15.], who designed, developed and evaluated a computer simulation to teach undergraduate students the scientific principle of competitive specialisation. Two interactive visualisations were designed: iconic images of ants and fruit (concrete) and dots and patches (abstract). By comparing learning gains between variations of type and order of presentation, i.e. abstract only, concrete only, Concreteness Fading (concrete followed by abstract), and concreteness introduction (abstract followed by concrete), the authors found that Concreteness Fading led to the best performance on both near and far transfer tests. However, these materials show a distinct interpretation of concrete and abstract representations. Both representations used in this study were situated similarly on the concrete to abstract spectrum, i.e. partially concrete to partially abstract may be a fair description of the sequence. With a relatively small progression along this spectrum, there may be less need for a three-step sequence to link mutual referents between the first and last representations explicitly.

Influenced by Bruner's Modes of Representation, Fyfe et al.'s [13.] Theoretical Model of Concreteness Fading provided several recommendations for designers inferred from the literature. The authors note the need for direct empirical validation. In Fyfe et al. [12.], the model was applied to the design of learning materials to teach children the principles of mathematical equivalence. The concrete representation combined metaphor (balance scales as an equation) and physical manipulation (placing objects on the scales to achieve balance). The abstract representation used mathematical notation only, and the intermediate representation showed an iconic representation of the balance scales annotated with its abstract referents. The materials were evaluated using the same experimental design as in Goldstone and Son [15.]. Concreteness Fading led to the highest performance on the transfer test, followed by abstract only, concreteness introduction, and finally, concrete only. Fyfe and colleagues found that Concreteness Fading was effective for all learners but presented evidence that it may be particularly effective for those with low existing knowledge in a domain [12.]. Trory et al. [28.] implemented the principles of the Theoretical Model of Concreteness Fading in a learning environment prototype to teach simple computing concepts to children. Applying the same experimental design, the authors found an identical ordering of groups to Fyfe et al., increasing the body of evidence for the premise that Concreteness Fading is more effective than all other variations of form and sequence of equivalent representations.

There are some examples of three-step progressions reporting successful results within educational technology design, where the concrete representation is implemented virtually rather than physically. This includes Koedinger and Anderson [18.], who used Concreteness Fading in their cognitive tutor system and found that presenting concrete algebraic word problems, then an intermediate representation, and then a symbolic representation brought higher learning gains. Bouck and colleagues [4., 5.] have also offered evidence that virtual manipulatives can be as effective or more effective than physical manipulatives in the Concreteness Fading progression.

Beyond Concreteness Fading research, embodied cognition theorists often argue for the learning benefits of physical interaction (e.g. [22.]), and Computer Science Unplugged-style activities have proven popular with teachers, with many believing that physical engagement is conducive to learning [21.]. In mathematics, a meta-analysis showed that physical manipulatives may be superior to abstract representations alone [9.]. However, when directly comparing the efficacy of physical and virtual manipulatives for learning, many studies have shown no significant difference in a range of domains, including chemistry [24.], physics [30.], and engineering [17.]. In fact, a meta-analysis of comparison studies showed a small effect in favour of virtual manipulatives [20.].

There appears to be no previously published empirical work on the effect of physicality within a Concreteness Fading sequence. We cannot reject the hypothesis that physicality may provide value in this context, but we should consider the potentially beneficial implications of a fully virtual design that enables all representations to be delivered within a single learning environment. The current body of empirical work on Concreteness Fading is limited, but established learning science theory may provide insight into directions for design and future research. For example, the split-attention principle refers to the extraneous cognitive load created within learners when mentally integrating multiple sources of information [2.], such as text and a diagram. Evidence shows that spatial congruity increases learning outcomes compared to spatial incongruity, for example, explaining a diagram through annotation is more effective than through an adjacent paragraph [26.]. A fully virtual implementation of a Concreteness Fading sequence enables mutual referents to remain in-place as they fade from concrete to abstract or, at least, gives the designer more control over their positioning when compared to distinct physical representations (physical manipulative to paper).

Evidence also shows that temporal congruity plays a role in finding and integrating mutual referents [19.], with the authors concluding that simultaneous presentation of words and pictures is more effective for learning than sequential presentation. This is important for designers of Concreteness Fading materials to note, given that it contradicts a fundamental premise. However, the learning materials referred to were designed such that the integration of mutual referents was necessary before understanding the concept, for example, text followed by a diagram. Within the Concreteness Fading approach, each representation contains all the information needed to understand each concept, with changes between representations limited to the progression from concrete to abstract. It may be reasonable to infer, however, that the observed effect of temporal congruity on cognitive load and learning outcomes may also impact a learner's ability to identify mutual referents within Concreteness Fading. If so, then the increased efficiency of representational change within a fully virtual implementation of the sequence could prove advantageous, particularly when compared to an intricate physical manipulative.

A fully virtual Concreteness Fading sequence may provide the opportunity to simplify the critical process of identifying mutual referents between representations by providing increased spatial and temporal contiguity. The increased flexibility of virtual design and presentation may provide further improvements. As previously stated, Concreteness Fading is commonly implemented as a two-step or three-step sequence. However, within a virtual implementation, the process of fading could be presented more gradually. Increasing the number of intermediate steps, with each step thus requiring a smaller representational change, may further scaffold a learner's progression from concrete to abstract.

3 AUGMENTED REALITY LEARNING ENVIRONMENT

3.1 Overview

This paper describes the third experiment in a project to investigate Concreteness Fading applied to teaching computing concepts to children. For experiment one [28.], a physical prototype of the learning environment was built to investigate how the form and sequence of representations affect gains for children learning introductory computer networking concepts. Networks were selected as a topic from the English National Curriculum Key Stage 2 (ages 7 to 11) Programme of Study [11.] for two reasons: 1) although deemed age-appropriate in some form, children would likely have low prior knowledge of abstract representations, i.e. edge lists and routing tables, and 2) networks can be effectively represented at multiple diverse positions on the concrete to abstract spectrum. The learning environment's design was inspired by the TMCF and the Computer Science Unplugged project [10.]. The three-step Concreteness Fading sequence began with a physical manipulative that implemented an engaging and intuitive conceptual metaphor. The intermediate representation presented an iconic visualisation of the manipulative, with progress towards the abstract in the positioning and description of conceptual features. The abstract representation showed an edge list and collection of routing tables. The experiment showed that Concreteness Fading improved upon concrete-only materials through increased learning gains in a pre-test / post-test design. For experiment two [27.], the physical representation was re-implemented as a tablet-based augmented reality equivalent to investigate whether and how the physicality of the concrete representation affects learning gains. The intermediate and abstract representations remained unchanged. A direct comparison found no significant difference in learning gains between the 'physical concrete' and 'virtual concrete' groups in a pre-test / post-test design. These findings did not support the TMCF position that the concrete representation should be physical.

For the current experiment, the 'virtual concrete' representation was extended to include all representations in a five-step Concreteness Fading sequence (while supporting two and three-step versions for comparison). The previous intermediate representation was updated, and two further intermediate representations were

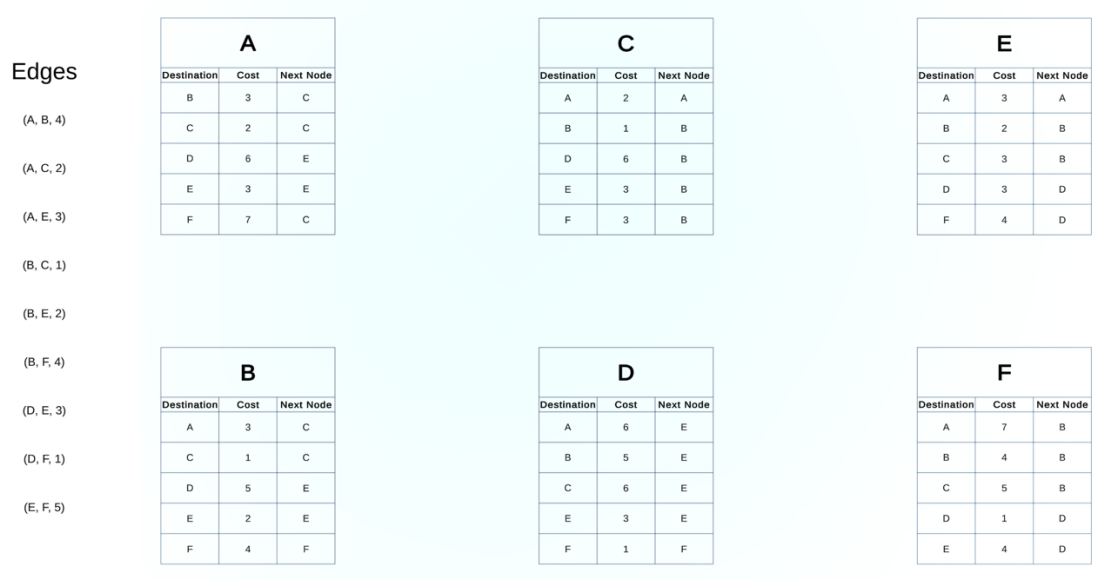

Figure 1: Abstract Representation in AR Learning Environment

added to allow for a more gradual fade. A narration script was written, incorporating the progression from concrete to abstract language, and delivered by the researcher. In future versions of the AR prototype, the narration could be provided within the app.

3.2 Abstract Representation

The materials aimed to teach children to effectively use an abstract representation of a network, shown as an edge list and simple routing tables (see Figure 1). The network contained six nodes, named A to F. Each node was connected to between two and four other nodes by edges, displayed only within an edge list on the left of the representation. Each edge was presented in the following format: "(node name, node name, cost)", for example, "(A, B, 4)". Each node had a corresponding routing table with three columns: "Destination", "Cost", and "Next Node", displayed on the right of the representation. The positioning of tables did not imply edges; for example, nodes B and E share an edge, but their tables are not adjacent. The script explained that the direct route between two nodes (the edge) did not guarantee the lowest cost route between them and that the routing tables must be consulted to locate this information.

3.3 Concrete Representation

The concrete representation was designed to introduce the network and routing concepts using a conceptual metaphor based on pirates' islands (see Figure 2). A group of six islands (nodes) are shown to the learner, each with a unique name, for example, "Pirates' Island" or "Shipwreck Bay". Each island was connected to between two and four other islands by white lines (edges). Lines were described as flight paths over which carrier parrots may fly to deliver messages on behalf of the residents of the islands. Each flight path was labelled with the time required to complete the flight (cost), for example, "2 mins". The flight path length did not imply the time required to travel over it. A set of direction signs (routing table) was positioned on each island. Each sign pointed along a flight path (next node) and displayed the name of an island (destination) and the total time required to reach that island (route cost). The script explained that following the direction signs guaranteed the fastest possible route to a destination, and parrots always wished to deliver messages as quickly as possible.

3.4 Intermediate Representations

A progression of three intermediate representations was designed to link the mutual referents between the concrete and abstract representations (see Table 1).

The Concrete B representation retained the visual and metaphorical details of Concrete A (previously described). However, the language used on all labels progressed towards the abstract; for example, "Shipwreck Bay" became "B" (island banner), "Musket Hill: 3 mins" became "C: 3" (direction sign), and "4 mins" became "(B, C, 4)" (flight path). This update to the flight path labels was also a conceptual progression due to including the names of the two connected islands.

The Intermediate A representation was a 2-dimensional image projected to the same size, shape, and position as the concrete representations. Islands were replaced by routing tables with two columns: "sign" and "points towards", and contained iconic images of the direction signs and the island banners. The format of the labelling was unchanged from Concrete B. Within each row, from left to right, the learner could read the name of a destination node, the route cost to the destination, and the next node along the route, matching the order of information in the abstract representation (previously described). The procedure for using the table was the same as that for the direction signs of the concrete representations: find the direction sign with the name of the required destination and travel to where it points. The visual representation of edges was

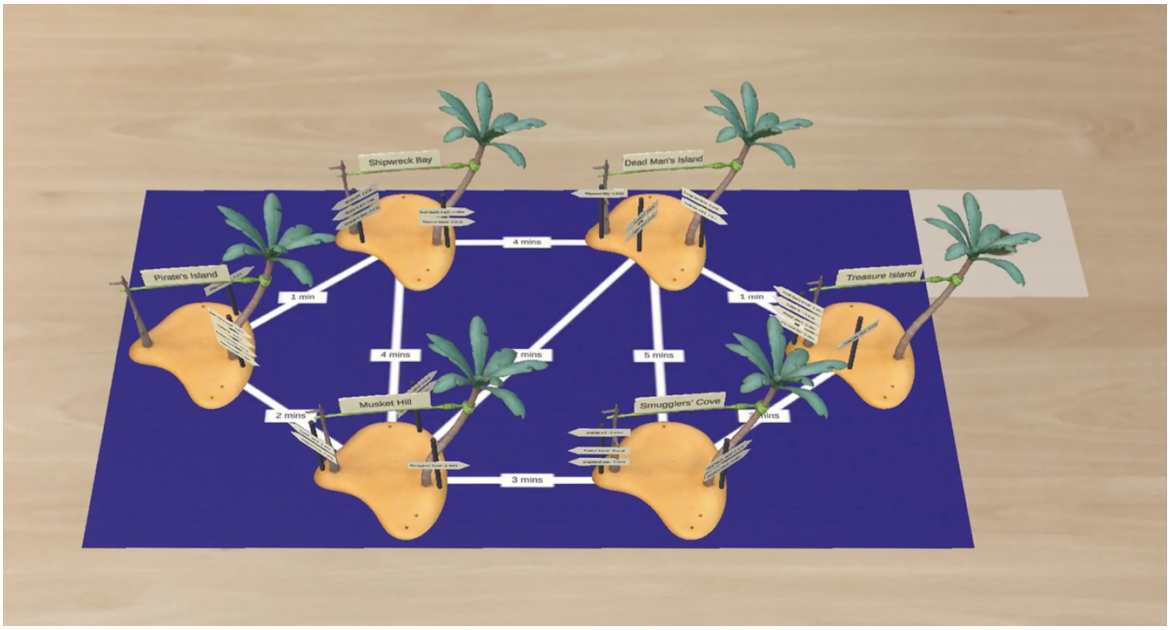

Figure 2: Concrete Representation in AR Learning Environment

Table 1: Representational Detail in Five-Step Fading

	Stage 1		Stage 2		Stage 3
	Concrete A	**Concrete B**	**Intermediate A**	**Intermediate B**	**Abstract**
	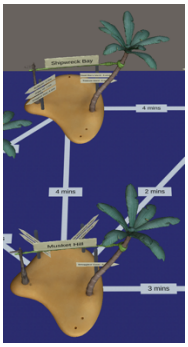	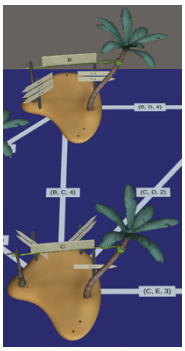	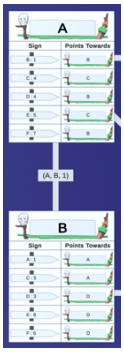	Edges (A, B, 1) 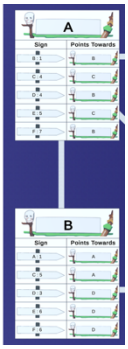	Edges (A, B, 4) 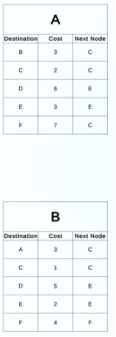
Language	Fully concrete	Partially abstract	Fully abstract	Fully abstract	Fully abstract
Edges	Fully concrete	Partially abstract	Partially abstract	Fully abstract + partially abstract	Fully abstract
Tables	Fully concrete	Fully concrete	Partially abstract	Partially abstract	Fully abstract

unchanged from Concrete B. However, in the script, the language used to describe the elements progressed to the abstract: islands were now nodes, flight paths were now edges, flight times were now costs, and the term "network" was introduced.

In the Intermediate B representation, routing tables were unchanged from the Intermediate A representation, but the edges progressed towards the abstract. The white lines that visually linked the nodes remained in place, but the labels were repositioned to a new area to the left of the routing tables and displayed in their final abstract form. In the script, the learner would first be directed to the white lines and then immediately to the edge list to read the relevant label.

3.5 Implementation of AR Learning Environment

The islands were first constructed using Playmobil® Pirates toys, wood, and cardboard (see [28.]) and then captured as 3D models using photogrammetry. The AR learning environment was developed using the Unity AR Foundation framework and installed on a 9.7" Apple iPad Pro. The iPad was enclosed within a rubber case

Table 2: Experimental Group Details

	Learning Stage 1	Learning Stage 2	Learning Stage 3
Two-step	Concrete A	Concrete A	Abstract
Three-step	Concrete A	Intermediate A	Abstract
Five-step	Concrete A / Concrete B	Intermediate A / Intermediate B	Abstract

with handles on either side (marketed for children), and an iPad stand was used intermittently. When viewing the iPad screen, the representations would appear to be positioned on the desk directly in front of the participant, at a size of approximately 100cm by 60cm. All five representations were contained within the learning environment, with the researcher manually switching between them using semi-transparent back and forward buttons positioned on the right side. The researcher and participants moved the iPad to control the display of the representation, for example, zooming to read labels or to focus on only the relevant elements. A Playmobil®parrot model was displayed in the concrete representations and could be moved between islands using the touchscreen. In the intermediate representations, the parrot was replaced by highlighting the node's name within its routing table. In both cases, the aim was to provide a simple cue to enable participants to keep track of their progress during a task. This support was not provided in the abstract representation.

4 METHOD

4.1 Design

This experiment used a between-groups design, with representational progression as the independent variable. Three groups were compared: two-step, three-step, and five-step. The initial concrete and final abstract representations were identical for all participants. Groups differed only in their use of intermediate representations between the concrete and abstract representations. All participants progressed through three stages of learning, as shown in Table 2, with each having an approximately equal duration. In the five-step group, the representation changed mid-way through stages one and two without significantly increasing the duration of the stage. All participants completed pre- and post-tests, with the difference being recorded as the dependent variable (learning gain).

4.2 Participants and Settings

Experimental sessions were conducted in a UK primary school (a state-maintained school for children aged between 5 and 11). The pre-tests were completed in the participants' regular classrooms. The learning activities and post-test were completed in a nearby corridor workspace. Distractions to participants caused by normal use of the corridor could not be prevented but mostly occurred at routine times. Session start times were balanced between groups to remove potential bias. All pupils in year five (ages 9 to 10) of the school were invited to participate. Information and consent sheets were distributed by teachers to pupils, who then sought written consent from their parents or carers. Ten pupils participated in an initial pilot (following which the learning activities were improved) and were not included in the main experiment. One participant withdrew from the experiment before completing the learning activities and post-test. The final sample size was 59, comprising 30 females and 29 males, aged between nine and ten years ($M = 9.93$, $SD = 0.30$). One participant was visually impaired, which required increasing the size of all materials by 100% in accordance with their teacher's usual practice. Participants were assigned to groups using blocked randomisation to achieve balance with the relatively small sample size. Participants were ordered based on the mean of their scores from recent assessments conducted by the school. Blocks of three were then defined and numbered consecutively, with 'block 1' having the lowest mean. Participants were randomly assigned to groups (see Table 3), so each block contained exactly one participant assigned to a two-step, a three-step, and a five-step group.

4.3 Materials

4.3.1 Learning Stages. All participants completed three stages of learning, each comprising four substages: narration > activity > narration > activity. The structure of the narration scripts and activities was closely matched across stages and between groups. However, the language used to describe concepts changed according to the representation; for example, islands became nodes, etc. In the five-step group, the representation changed at the mid-point of both stage one and stage two (immediately following the first activity). Therefore, additional lines were included in the narration script to explain these changes, but the increase in overall duration was minimal. Table 4 shows excerpts of the narration script for each group, taken from the second learning stage. The activities comprised seven unique tasks, phrased to closely match the pre- and post-test questions but using language appropriate for the current representation. For example, when using a concrete representation, "How many flight paths connect Shipwreck Bay directly to other Islands?" would, in abstract form, be phrased as "How many edges connect Node A directly to other Nodes?". Finding the solution in the concrete representation required counting the flight paths visibly connected to the island, whereas in the abstract representation, the participant must carefully read each row of the edge list, counting only those edges that contain the letter A.

4.3.2 Pre- and Post-Test. A near-transfer test was designed to measure participants' ability to interpret and use a partially abstract representation of a network. During the learning stages, no context was provided for the abstract representation of the network. However, the test described the network as "a fictional example of how a computer programme might represent the data lines and routing tables that make up the internet". Nodes were described as routers, edges as data lines, and cost as (transmission) time in milliseconds. This terminology was introduced through an introduction sheet that was read aloud by the researcher and phrased in such a way as

Table 3: Gender Distribution Between Groups

	Females	Males	Total
Two-Step	9	11	20
Three-step	10	9	19
Five-step	11	9	20
N	30	29	59

Table 4: Script Extract Comparison Across Conditions (Learning Stage 2)

Two-step: Concrete B	Three-step: Intermediate A	Five-step: Intermediate B
		I have moved the labels from the edges and put them on the left, in a list. We'll call this the edge list. [researcher focuses the representation on the edge list]
Just like before, the parrot always wants to deliver the message as quickly as possible.	Just like before, we always want to have the lowest cost route across the network.	Just like before, we always want to have the lowest cost route across the network.
The parrot is on Pirate's Island, and it wants to deliver a message to a friend on Dead Man's Island. [researcher focuses the representation on the two islands and connecting flight path] There is a flight path that we can see, directly connecting the two, so the parrot can fly straight there.	We're at node A, and we want to go to node D. [researcher focuses the representation on the two tables and connecting line] There is an edge that we can see, directly connecting the two, so we can go straight there.	We're at node A, and we want to go to node D. [researcher focuses the representation on the two tables and connecting line] There is an edge that we can see, directly connecting the two, so we can go straight there.
We can see that the label on the flight path is "5 mins", so the direct route from Pirate's Island to Dead Man's Island would take the parrot five minutes to fly across.	We can see that the label on the edge is (A, D, 5), so the direct route from A to D would cost 5.	We don't know how much that costs because we can't see the label anymore, so let's look it up in the edge list. [researcher focuses the representation on the edge list] There we can see (A, D, 5), so the direct route from A to D would cost 5.

to encourage the participants to connect concepts to the previous abstract representation, but without explicit linking. Routers were named according to the cities they resided in, for example, "Boston" and "San Francisco". All cities were in the USA, rather than the UK, to minimise the participants' ability to infer proximity and thus relative data transmission times. The complexity of the test network was greater than that of the learning network, containing nine routers and 15 data lines. The design of the representation was identical to the previous abstract representation: an edge list on the left of the page and routing tables to the right. Two versions of the representation were created, A and B, differing only in the names of cities and ordering of edges and routing tables. Participants were given one as the pre-test and the other as the post-test.

The test comprised eight multiple-choice questions, with four focused primarily on the edge list and four on the routing tables. The latter was not explicitly stated to participants, but most questions contained the words "directly" or "fastest" as a cue. For example, Q3 "How long does it take for data to travel between Phoenix and Houston directly?" and Q5 "What is the fastest that data could be sent from Boston to Miami?" In both cases, the MCQ responses were designed such that a seemingly plausible answer could be derived from either the edge list or the routing tables. After correctly identifying what the question referred to, the participant would search for and interpret the information.

4.4 Procedure

The experiment ran from November 2022 to February 2023 and was conducted with three classes concurrently. Each participant completed two sessions: the collectively administered pre-test session (S1) and the individually administered learning and post-test session (S2). The mean period between S1 and S2 was 77.46 days (SD = 24.63, X(1) = 35, X(n) = 112). Two S2 sessions were run each day at approximately 1:15 p.m. and 2:10 p.m., with the participant for each S2 session chosen pseudo-randomly by their classroom teacher immediately before starting.

The S1 sessions lasted approximately 15 minutes. After an introduction from the researcher, the pre-test was distributed to the participants face-down and alternating between edition A and edition B, such that no participant was seated next to another with the same edition of the pre-test. Participants were asked to turn over

Table 5: Learning Gains Summary - Descriptive

	Pre-test	SD	Post-test	SD	Gain	SD
Overall	0.81	0.88	3.9	1.55	3.09	1.52
Females	0.77	0.83	4.15	1.49	3.38	1.45
Males	0.86	0.93	3.64	1.58	2.78	1.54
Upper Blocks	0.88	0.9	4.15	1.66	3.27	1.57
Lower Blocks	0.74	0.86	3.64	1.4	2.9	1.47

the introductory page and follow along while the researcher read aloud from it. Participants were then asked to turn over the remaining pages and given seven minutes to complete the multiple-choice worksheet. Finally, participants were asked to put their pencils down, and the pre-tests were collected.

The S2 sessions lasted approximately 50 minutes. Participants were first given an overview of the session and asked for verbal assent to continue, with all agreeing to do so. The procedure for each of the three learning stages was the same. The researcher began by revealing the representation to the participant and then narrated while moving the iPad to display the relevant features. It was planned that if a participant were to query a point during narration, the researcher would repeat the relevant portion of the script a maximum of two times before moving on. However, no participant needed more than one repetition.

Next, the researcher placed the activity sheet on the desk and read the task aloud while directing the participants' attention to the sheet. Participants were then asked to take control of the iPad and explore the representation to find the answer. Answers were given verbally to the researcher, who noted the response on the sheet. If a participant was unable to understand or answer the question, the researcher used his judgment to repeat portions of the script that may support the task. If the participant was still unable to complete the question, the correct solution was demonstrated. After the learning stages were finished, the post-test was completed as in the procedure for S1 described above.

5 RESULTS

Learning gains were calculated by subtracting pre-test scores from post-test scores (see Table 5). The overall mean learning gain score was 3.09 (SD = 1.52). An independent-samples t-test was run to determine if gender differences existed in learning gains. Females (M = 3.38, SD = 1.45) achieved higher learning gains than males (M = 2.78, SD = 1.54), but the difference was not statistically significant (M = 0.61, 95% CI [-0.17, 1.39], $t(57)$ = 1.56, p = 0.063). Linear regression showed that assignment block number (based on past scores on relevant assessments) could statistically significantly predict learning gains (R^2 = 0.086, $F(1,57)$ = 2.32, p = 0.024). Results for blocks were compared by defining "upper blocks" (higher scores on relevant past assessments) (n = 30) and "lower blocks" (n = 29), splitting the sample at the median.

A one-way ANOVA was conducted to determine the effect of intermediate representations on learning gains. There were no outliers, as assessed by boxplot; data were normally distributed for each group, as assessed by Shapiro-Wilk ($p > 0.05$); and there was homogeneity of variances, as assessed by Levene's test of homogeneity of variances (p = 0.67). Learning gains were statistically significantly different between groups ($F(2,56)$ = 3.67, p = 0.032, η^2 = 0.12), increasing from the two-step (M = 2.38, SD = 1.46) to three-step (M = 3.37, SD = 1.55) to five-step (M = 3.53, SD = 1.34) group, in that order (see Table 6 and Figure 3). Planned contrasts were conducted to test hypotheses one and two. The mean increase in learning gains between the two-step and three-step groups (hypothesis one) was statistically significant (M = 0.99, 95% CI, 0.61 to 0.99, p = 0.037). The mean increase in learning gains between the three-step and five-step groups (hypothesis two) was not statistically significant (M = 0.16, 95% CI, -0.78 to 1.09, p = 0.738).

For the upper blocks, learning gains were not statistically significantly different between groups ($F(2,27)$ = 2.36, p = 0.397) increasing from the two-step (M = 2.85, SD = 1.31) to three-step (M = 3.15, SD = 1.73) to five-step (M = 3.80, SD = 1.64) group, in that order. For the lower blocks, learning gains were statistically significantly different between groups ($F(2,26)$ = 4.62, p = 0.019, η^2 = 0.26), increasing from the two-step (M = 1.90, SD = 1.51) to five-step (M = 3.25, SD = 0.98) to three-step (M = 3.61, SD = 1.39) group, in that order (see Figure 4). Planned contrasts were conducted. The mean increase in learning gains between the two-step and three-step groups was statistically significant (M = 1.71, 95% CI, 0.48 to 2.95, p = 0.008). The mean increase in learning gains between the three-step and five-step groups was not statistically significant (M = -0.36, 95% CI, 0.87 to -1.60, p = 0.553). An independent-samples t-test was run to examine the difference in learning gains between the lower and upper blocks within the two-step group only but was found not to be statistically significant (M = 0.95, 95% CI [-0.38, 2.28], $t(18)$ = 1.50, p = 0.075).

6 DISCUSSION

This experiment investigated the explicit linking of mutual referents between representations in a Concreteness Fading sequence implemented in augmented reality. Three approaches to linking were designed (two-step, three-step, and five-step), and two research questions were examined. There is some evidence that the Concreteness Fading approach may be particularly effective for those with low existing knowledge in a domain [12.]. The design of this experiment controlled for this potentially confounding variable by using blocked randomisation based on each participant's mean of recent assessment scores provided by the school. Splitting the sample at the median block, we also examined each research question when applied to only the upper blocks (participants with means above the sample median) and the lower blocks (participants with means below the sample median).

Table 6: Learning Gains Across Groups – Descriptive

	Pre-test	SD	Post-test	SD	Gain	SD
Overall	0.81	0.88	3.9	1.55	3.09	1.52
Two-step	0.83	0.96	3.2	1.63	2.38	1.46
Three-step	0.82	0.82	4.18	1.22	3.37	1.55
Five-step	0.8	0.88	4.33	1.57	3.53	1.34

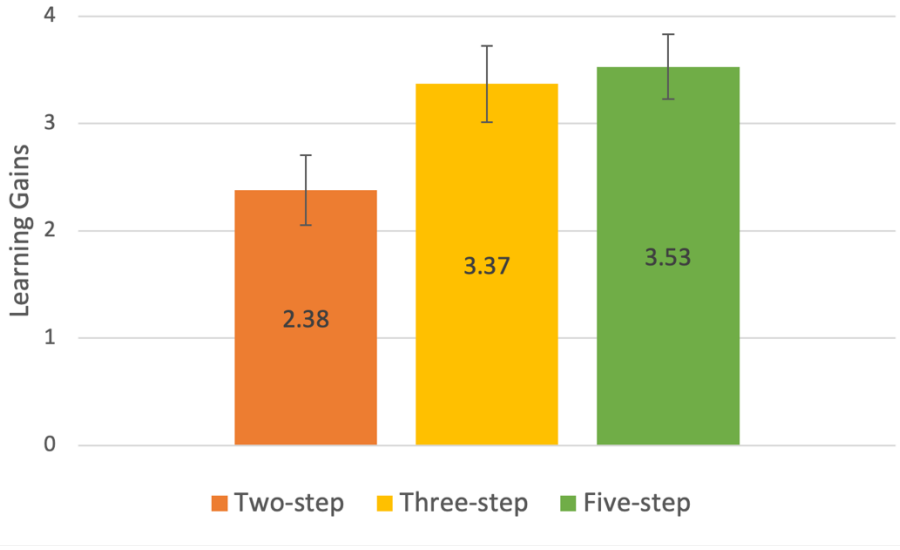

Figure 3: Learning Gains by Group

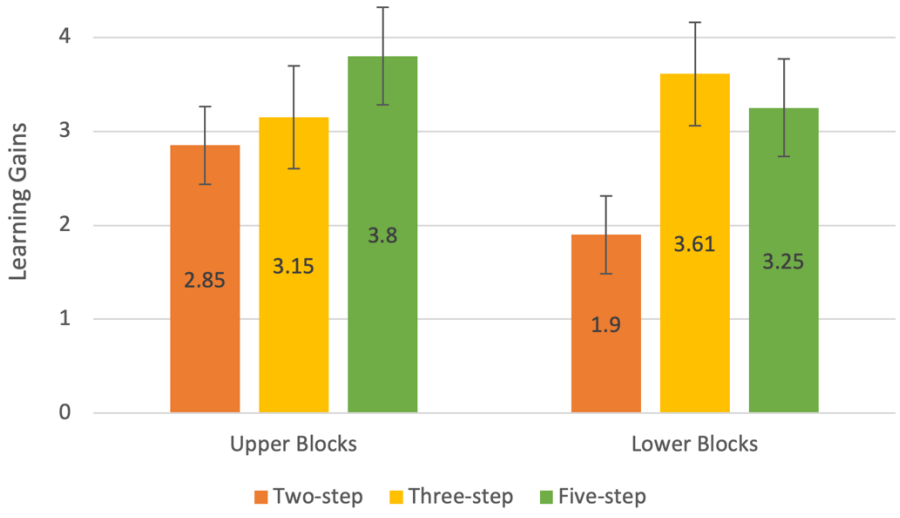

Figure 4: Learning Gains by Group - Upper and Lower Blocks

A one-way ANOVA was conducted, including all three groups, to confirm the fundamental premise of this study: the approach to linking concrete and abstract representations within Concreteness Fading influences learning outcomes. Learning gains were found to be statistically significantly different between groups. In the lower blocks, between-groups differences were found to be statistically significant, with an increase in effect size compared to the full sample. In the upper blocks, learning gains were not statistically significantly different between groups, but the ordering of differences did match the full sample. A comparison of learning gains between the two-step groups in the upper and lower blocks was not statistically significant. Considering the relatively small subsample ($n = 20$) of this statistical test, this outcome is unsurprising and further research is required to determine validity. However, if confirmed, it would imply that learners with lower previous assessment scores may struggle to identify and link mutual referents between a sequentially presented concrete and abstract representation compared to learners with higher previous assessment scores.

In the first of two research questions, we examined how linking the concrete and abstract representations through a single intermediate representation affects learning gains (RQ1). Two groups were compared, two-step and three-step, to evaluate the hypothesis that the three-step group would have higher learning gains than the two-step group. Learning gains increased from the two-step group to the three-step group, and the difference was statistically significant. These results confirm our hypothesis. In the lower blocks, learning gains increased from the two-step group to the three-step group, and the difference was statistically significant. In the upper blocks, learning gains increased from the two-step group to the three-step group, but the difference was not statistically significant. We cannot infer the efficacy of Concreteness Fading from this experimental design, but these results do show that for students with lower previous assessment scores, explicitly linking the mutual referents between concrete and abstract representations increases learning gains. Further research is needed to confirm if this may also be true for the students with higher previous assessment scores. A larger sample size would likely be needed, as these results imply that if an effect exists in this group, it would be smaller.

In the second of two research questions, we examined whether a more gradual fade through an increased number of intermediate representations affects learning gains (RQ2). Two groups were compared, three-step and five-step, to evaluate the hypothesis that the five-step group would have higher learning gains than the three-step group. Learning gains increased from the three-step group to the five-step group, but this difference was not statistically significant, and we must reject our hypothesis. In the upper blocks, there was a trend for learning gains to increase from the three-step group to the five-step group, as predicted. Conversely, in the lower blocks, learning gains *decreased* from the three-step group to the five-step group. In the context of previously finding that adding an explicit linking representation increased learning gains for the lower blocks, this result was unexpected. However, considering the lack of statistical significance and potential influence of noise in low sample sizes (n = 19), it is most likely that this result is not meaningful. To be cautious, it may be of value to investigate whether an increased number of intermediate representations may be detrimental to learning for students with lower previous assessment scores, potentially due to less time to familiarise themselves with a new representation before it changes.

6.1 Limitations and Future Work

To maximise ecological validity, we ran our experiment in an authentic setting: all data were collected in either a classroom or a shared corridor workspace in a primary school. It was planned that learning gains, measured quantitatively via pre- and post-tests, would be the primary data for analysis. However, we also planned to collect audio/visual data on the individual learning sessions to enrich our evaluations with further metrics and qualitative data. Unfortunately, this proved impractical due to the nature of the space. The audio/visual equipment could not be positioned to exclude the recording of non-participants and created a potential trip hazard in a busy corridor. Furthermore, to reduce the burden placed on the school, it was decided that each participant would complete only one post-test, presented immediately after their learning session. These data provide a measure of understanding, but we cannot infer participants' ability to recall after an extended period. An additional delayed post-test may have given insight into whether explicitly linking mutual referents provided increased value over time. For further discussion of challenges and potential limitations faced during in-school evaluations of educational technology, see [29.].

For the purposes of this experiment, the learning stages and representations were designed to enable the completion of a full Concreteness Fading sequence within a single 50-minute session. However, in some implementations of Concreteness Fading, fading occurs over multiple sessions (within-lesson vs across-lesson fading [25.]), allowing the learner more time to manipulate representations and, potentially, better understand the relationship between them. How the duration or number of sessions may affect the required number of intermediate linking representations is unknown, and a longitudinal study may be required to answer this question. Furthermore, the results of this experiment imply that the explicit linking of representations has a greater positive effect on the lower blocks than the higher blocks, i.e. some learners need more representational scaffolding than others. Therefore, it may be that the optimal implementation of Concreteness Fading would require an adaptive system in which learners progress along the concrete-to-abstract spectrum at a pace uniquely appropriate for them rather than a generic pre-determined point in the session (as in this experiment). This should be investigated further in future work.

The complexity of the concepts being learned and the extent of the differences between the concrete and abstract representations may also have a crucial role in determining the number of intermediate steps required to maximise learning gains. For the AR prototype and learning context described in this paper, the evidence suggests the three-step progression is as effective as the five-step. However, future work should investigate whether this finding holds across other domain areas and learning activities.

7 CONCLUSION

This paper has concentrated on the approach to representational fading within the Concreteness Fading paradigm and examined it within the context of teaching computing concepts to children aged 9-10. In particular, it has explored the effect on learning gains of distributing a representational fade across zero, one, or three intermediate representations while maintaining constant overall time on task. To conduct this study, an augmented reality learning environment was implemented on an iPad to enable the children to view and interact with the different forms of representation. Results indicate that Concreteness Fading may be more effective when an intermediate representation explicitly links the mutual referents between the concrete and abstract representation. However, a more gradual fade across multiple intermediate representations may not significantly increase learning gains any further. A secondary and not unexpected result was that Concreteness Fading seemed to play a more helpful role for children with lower mean scores on previous assessments compared with children who had scored more highly.

8 SELECTION AND PARTICIPATION OF CHILDREN

Ethical approval for this study was obtained from the University of Sussex Sciences & Technology Cross-Schools Research Ethics Committee (SCITEC C-REC). All pupils in year five (approx. 100 children, aged 9-10) of a local primary school were invited to participate. The informed consent procedure involved written parental permission for all participants to take part in the tests and learning activities. All child participants also gave verbal assent in response to an assent script prior to taking part in research activities. The final sample size (N=59) and comprised 30 girls and 29 boys.

ACKNOWLEDGMENTS

Many thanks to the teachers and children who donated their time and effort to our work. This study is part of doctoral research funded by the Engineering and Physical Sciences Research Council (EPSRC).

REFERENCES

[1.] Shaaron Ainsworth. 1999. The functions of multiple representations. *Computers and Education* 33, 2–3: 131–152.

[2.] Paul Ayres and John Sweller. 2014. The split-attention principle in multimedia learning. In *The Cambridge Handbook of Multimedia Learning* (2nd ed.), Richard E. Mayer (ed.). Cambridge University Press, 206–226.

[3.] Tim Bell and Jan Vahrenhold. 2018. CS unplugged—How is it used, and does it work? In Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). https://doi.org/10.1007/978-3-319-98355-4_29

[4.] Emily C. Bouck, Laura Bassette, Jordan Shurr, Jiyoon Park, Jackie Kerr, and Abbie Whorley. 2017. Teaching Equivalent Fractions to Secondary Students With Disabilities via the Virtual–Representational–Abstract Instructional Sequence. *Journal of Special Education Technology* 32, 4. https://doi.org/10.1177/0162643417727291

[5.] Emily C. Bouck, Rajiv Satsangi, Teresa Taber Doughty, and William T. Courtney. 2014. Virtual and concrete manipulatives: A comparison of approaches for solving mathematics problems for students with autism spectrum disorder. *Journal of Autism and Developmental Disorders* 44, 1. https://doi.org/10.1007/s10803-013-1863-2

[6.] Jerome S. Bruner. 1966. *Toward a theory of instruction.* Harvard University Press, Cambridge, MA.

[7.] Frances M. Butler, Susan P. Miller, Kevin Crehan, Beatrice Babbitt, and Thomas Pierce. 2003. Fraction Instruction for Students with Mathematics Disabilities: Comparing Two Teaching Sequences. *Learning Disabilities Research & Practice* 18, 2. https://doi.org/10.1111/1540-5826.00066

[8.] Elisa Nadire Caeli and Aman Yadav. 2020. Unplugged Approaches to Computational Thinking: a Historical Perspective. *TechTrends* 64, 1. https://doi.org/10.1007/s11528-019-00410-5

[9.] Kira J. Carbonneau, Scott C. Marley, and James P. Selig. 2013. A meta-analysis of the efficacy of teaching mathematics with concrete manipulatives. *Journal of Educational Psychology* 105, 2. https://doi.org/10.1037/a0031084

[10.] CS Unplugged. Principles. Retrieved April 14, 2024 from https://www.csunplugged.org/en/principles/

[11.] Department for Education. 2013. *Computing programmes of study: key stage 1 and 2. National curriculum in England.* Retrieved April 14, 2024 from https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/239033/PRIMARY_national_curriculum_-_Computing.pdf

[12.] Emily R. Fyfe, Nicole M. McNeil, and Stephanie Borjas. 2015. Benefits of "concreteness fading" for children's mathematics understanding. *Learning and Instruction* 35: 104–120.

[13.] Emily R. Fyfe, Nicole M. McNeil, Ji Y. Son, and Robert L. Goldstone. 2014. Concreteness fading in mathematics and science instruction: A systematic Review. Educational Psychology Review 26, 1: 9–25.

[14.] Emily R. Fyfe and Mitchell J. Nathan. 2019. Making "concreteness fading" more concrete as a theory of instruction for promoting transfer. *Educational Review* 71, 4. https://doi.org/10.1080/00131911.2018.1424116

[15.] Robert L. Goldstone and Ji Y. Son. 2005. The transfer of scientific principles using concrete and idealized simulations. *Journal of the Learning Sciences* 14, 1: 69–110.

[16.] Amy M. Johnson, Jana Reisslein, and Martin Reisslein. 2014. Representation sequencing in computer-based engineering education. Computers and Education 72. https://doi.org/10.1016/j.compedu.2013.11.010

[17.] David Klahr, Lara M. Triona, and Cameron Williams. 2007. Hands on what? The relative effectiveness of physical versus virtual materials in an engineering design project by middle school children. *Journal of Research in Science Teaching* 44, 1. https://doi.org/10.1002/tea.20152

[18.] Kenneth R. Koedinger and John R. Anderson. 1998. Illustrating Principled Design: The Early Evolution of a Cognitive Tutor for Algebra Symbolization. *Interactive Learning Environments* 5, 1. https://doi.org/10.1080/1049482980050111

[19.] Richard E. Mayer and Richard B. Anderson. 1992. The Instructive Animation: Helping Students Build Connections Between Words and Pictures in Multimedia Learning. *Journal of Educational Psychology* 84, 4. https://doi.org/10.1037/0022-0663.84.4.444

[20.] Patricia S. Moyer-Packenham and Arla Westenskow. 2013. Effects of virtual manipulatives on student achievement and mathematics learning. *International Journal of Virtual and Personal Learning Environments* 4, 3. https://doi.org/10.4018/jvple.2013070103

[21.] Sue Sentance and Andrew Csizmadia. 2017. Computing in the curriculum: Challenges and strategies from a teacher's perspective. *Education and Information Technologies* 22, 2: 469–495.

[22.] Lawrence Shapiro and Steven A. Stolz. 2019. Embodied cognition and its significance for education. *Theory and Research in Education* 17, 1. https://doi.org/10.1177/1477878518822149

[23.] Evelyn Sowell. 1974. Another Look at Materials in Elementary School Mathematics. *School Science and Mathematics* 74, 3. https://doi.org/10.1111/j.1949-8594.1974.tb09231.x

[24.] Andrew T. Stull, Trevor Barrett, and Mary Hegarty. 2013. Usability of concrete and virtual models in chemistry instruction. *Computers in Human Behavior* 29, 6. https://doi.org/10.1016/j.chb.2013.06.012

[25.] Sangho Suh, Martinet Lee, and Edith Law. 2020. How do we design for concreteness fading: Survey, general framework, and design dimensions. In *Proceedings of the Interaction Design and Children Conference, IDC 2020.* https://doi.org/10.1145/3392063.3394413

[26.] John Sweller, Paul Chandler, Paul Tierney, and Martin Cooper. 1990. Cognitive Load as a Factor in the Structuring of Technical Material. *Journal of Experimental Psychology: General* 119, 2. https://doi.org/10.1037/0096-3445.119.2.176

[27.] Anthony Trory. 2024. Concreteness Fading in Primary School Computing [Unpublished doctoral thesis]. University of Sussex, Brighton.

[28.] Anthony Trory, Kate Howland, and Judith Good. 2018. Designing for concreteness fading in primary computing. In *IDC 2018 - Proceedings of the 2018 ACM Conference on Interaction Design and Children.* https://doi.org/10.1145/3202185.3202748

[29.] Megan Venn-Wycherley, Ahmed Kharrufa, Susan Lechelt, Rebecca Nicholson, Kate Howland, Abrar Almjally, Anthony Trory, and Vidya Sarangapani. 2024. The Realities of Evaluating Educational Technology in School Settings. *ACM Transactions on Computer-Human Interaction* 31, 2. https://doi.org/10.1145/3635146

[30.] Zacharias C. Zacharia and Georgios Olympiou. 2011. Physical versus virtual manipulative experimentation in physics learning. *Learning and Instruction* 21, 3. https://doi.org/10.1016/j.learninstruc.2010.03.001