Using Machine Learning to Overcome the Expert Blind Spot for Perceptual Fluency Trainings

Martina A. Rau(✉), Ayon Sen, and Xiaojin Zhu

University of Wisconsin – Madison,
1025 W Johnson Street, Madison, WI 53706, USA
marau@wisc.edu

Abstract. Most STEM domains use multiple visual representations to illustrate complex concepts. While much research has focused on helping students make sense of visuals, students also have to become perceptually fluent at translating among visuals fast and effortlessly. Because perceptual fluency is acquired via implicit, nonverbal processes, perceptual fluency trainings provide simple classification tasks that vary visual features across numerous examples. Prior research shows that learning from such trainings is strongly affected by the sequence of the examples. Further, prior research shows that perceptual fluency trainings are most effective for high-performing students but may confuse low-performing students. We propose that a lack of benefits for low-performing students may result from a perceptual expert blind spot of instructors who typically develop perceptual fluency trainings: expert instructors may be unable to anticipate the needs of students who do not see meaningful information in the visuals. In prior work, we used a machine-learning approach to develop a sequence of example visuals of chemical molecules for low-performing students. This study tested the effectiveness of this sequence in comparison to an expert-generated sequence in a randomized experiment as part of an undergraduate chemistry course. We determined students' performance based on log data from an educational technology they used in the course. Results show that the machine-learned sequence was more effective for low-performing students. The expert sequence was more effective for high-performing students. Our results can inform the development of perceptual-fluency trainings for adaptive educational technologies.

Keywords: Multiple visuals · Perceptual fluency · Sequencing · Machine learning

1 Introduction

Visual representations are often used to illustrate concepts in science, technology, engineering, and math (STEM) instruction [1–3]. For example, chemistry instruction on bonding typically uses visuals such as Lewis structures and space-filling models of molecules (see Fig. 1) [4]. Multiple visual representations can enhance learning because they provide complementary information about the to-be-learned concepts [5–7]

S. Isotani et al. (Eds.): AIED 2019, LNAI 11625, pp. 406–418, 2019.
https://doi.org/10.1007/978-3-030-23204-7_34

(e.g., the Lewis structure shows how many electrons form bonds, the space-filling model shows the geometry of the molecule). However, multiple visual representations can impede learning if students are unable to make such connections among them [7, 8].

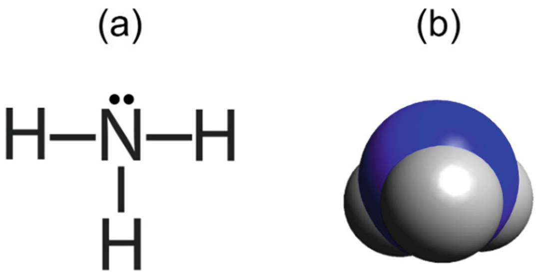

Fig. 1. (a) Lewis structure and (b) space-filling model of ammonia.

Most prior research on connection making among visuals has focused on helping students make sense of the connections by prompting them to explain differences and similarities between visuals [5, 9, 10]. For example, a student has to explain similarities in how the visuals in Fig. 1 show atoms: the H's in the Lewis structure in Fig. 1a correspond to the white spheres in the space-filling model in Fig. 1b because both show hydrogen atoms. They have to explain differences, for example that the dots in the Lewis structure show electrons that are not shown in the space-filling model.

By contrast, little research has focused on the role of *perceptual fluency* in students' learning: that is, the ability to quickly and effortlessly integrate information from multiple visuals [7, 11]. For example, students need to immediately see that the visuals in Fig. 1 show the same molecule and translate between them as fluently as bilinguals translate between languages [7, 12, 13]. Perceptual fluency frees cognitive resources for future learning and effortful conceptual reasoning [13, 14]. Because perceptual fluency is acquired via implicit, inductive, nonverbal processes [15, 16], instructional trainings enhance perceptual fluency by exposing students to a sequence of many simple problems that ask them to quickly judge what a visual shows [7, 11]. Perceptual-fluency training sequences make use of the contrasting cases principle, so that consecutive examples vary visual features so they draw students' attention to relevant features [7, 11]. However, while such trainings have proven effective for high-performing students, they are often ineffective for low-performing students [17, 18].

It is possible that the ineffectiveness of perceptual-fluency trainings for low-performing students is due to a perceptual expert blind spot on the part of the designers of instructional sequences. Sequences that present contrasting cases may be appropriate for high-performing students who already have a preliminary understanding of which visual features are meaningful. However, such sequences may be ineffective for low-performing students who have little prior knowledge about these features. For an instructional designer who is an expert in processing the visuals, it may be difficult to empathize with students who do not "see" meaning in the visuals [3, 7, 19].

In prior work [38], we used a machine-learning algorithm to develop a sequence of visuals for low-performing students. The machine-learned sequence was more effective than an expert sequence for participants from Amazon's Mechanical Turk (MTurk) service. However, it is unclear whether the machine-learned sequence is more effective than an expert sequence in a realistic learning context. To this end, we conducted an experiment with undergraduate students in a chemistry course.

2 Theoretical Background

In the following, we briefly review prior research on perceptual fluency as well as our own prior work on developing an instructional sequence of visual representations for students who lack prior knowledge about the visual representations.

2.1 Inductive Learning of Perceptual Fluency

In contrast to a large body of research on verbally mediated explanation-based sense making of visuals (see [5, 7] for overviews), research on the role of perceptual fluency in education is still relatively novel. This line of research builds on the expert-novice literature, which shows that experts see meaningful connections among visuals like the ones in Fig. 1 quickly and almost automatically [15, 16, 20, 21]. Experts are able to "at a glance" see meaning in visuals because translating and combining information from them takes little or no cognitive effort [11, 22]. This highlevel of efficiency at translating among visuals results from perceptual chunking: visual features of the representations serve to retrieve schemas that describe conceptual information from long-term memory [23, 24]. This high efficiency frees cognitive resources for higher-order thinking [12, 23] and is considered an important learning goal in many STEM domains [3].

Cognitive learning theories suggest that students acquire perceptual fluency via *inductive processes* involved in pattern learning [21, 25]. These processes involve both bottom-up mechanisms (e.g., a visual feature cues the retrieval of a conceptual schema) and top-down processes (e.g., conceptual schemas direct a student's visual attention to relevant visual features). Such inductive processes are considered to be non-verbal [23, 25] because verbal reasoning is not necessary [11, 13] and may even interfere with the acquisition of perceptual fluency [19, 26]. Consequently, students do not require direct instruction to become perceptually fluent, but rather acquire perceptual fluency through experience-based instructional sequences that expose them to many visuals [11, 21].

2.2 Perceptual-Fluency Trainings

In line with cognitive learning theories, perceptual-fluency trainings typically expose students to many examples, for instance in classification problems that ask students to quickly translate between visuals while providing simple feedback on whether the classification is correct or incorrect [27–30]. Because perceptual learning is strongly influenced by the order in which visuals are presented [7], perceptual-fluency trainings purposefully sequence visuals so that consecutive problems vary irrelevant visual features but repeatedly expose students to relevant features [11]. Through experience with many examples, students inductively learn to attend to relevant visual features [11].

The effectiveness of perceptual-fluency trainings has been demonstrated in many STEM domains, including math [31, 32] and chemistry [17, 33]. For example, Fig. 2 shows a perceptual-fluency problem from a training we developed for chemistry students. Each problem asks students to judge whether two visual representations show the same molecule or not. They receive many such problems in a row, sequenced according to the principles just described.

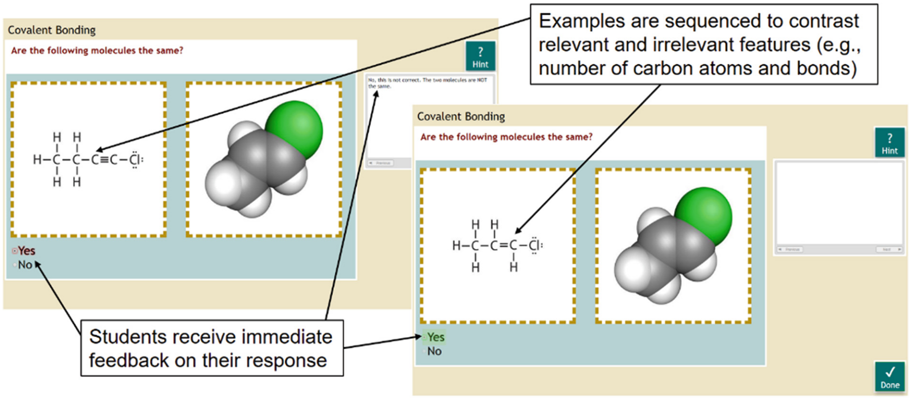

Fig. 2. Example perceptual-fluency problems in the expert-generated training sequence.

However, positive effects of such perceptual-fluency trainings have been confined to students who have substantial prior knowledge about the visual representations and the concepts they show [17, 18]. Indeed, much of the pioneering work on perceptual-fluency trainings in STEM was conducted with students after they had received considerable conceptual instruction and problem-solving practice with the visual representations [32, 34]. Further, recent research shows that students benefit from perceptual-fluency trainings only after they have acquired conceptual knowledge about connections among the visual representations [17, 18].

It is possible that students need to acquire conceptual knowledge about the visual representations before they can become perceptually fluent with them, as proposed by prior research [18]. An alternative explanation is that the lack of effectiveness of perceptual-fluency trainings for students with low prior knowledge may result from an expert blind spot on the part of instructional designers. The expert blind spot is a known phenomenon in the literature on conceptual learning because it can interfere with instructors' anticipation of student difficulties, which may hamper their ability to develop effective instruction [35]. However, we are not aware of any research on perceptual fluency in education that addressed this phenomenon. Specifically, while it is well documented that experts are unaware of why or how they perceive information in a certain way [19, 23, 36], knowledge of this lack of awareness has not informed the instructional design of perceptual-fluency trainings. For example, it is possible that instructors create sequences of visuals that inadvertently assume that students pay attention to specific visual features. In our prior work, which we review next, we used a machine-learning approach that is not subject to expert blind spot biases to create sequences for perceptual trainings for novice students.

2.3 Machine-Learned Sequences for Perceptual-Fluency Trainings

In prior work, we drew on Zhu's machine-teaching paradigm [37], which inverses typical machine-learning approaches to reverse-engineer optimal training sequences (for a detailed report see [38]). Given a cognitive model of a student learning to translate between pairs of visuals, the machine-teaching algorithm identifies a sequence of pairs that is most effective for training the cognitive model. To this end, the algorithm draws possible perceptual-fluency problems (e.g., Figure 2) from an underlying training distribution (not necessarily independently and identically distributed) to form a training sequence. Then, the cognitive model is trained with this sequence. Specifically, we used a feed-forward artificial neural network (ANN) as our learning algorithm. The inputs to the ANN were two feature vectors that corresponded to the visual features of the two visuals in a given perceptual-fluency problem. The ANN had mapped each of the two feature vectors to an embedding that corresponds to a space where visuals of similar molecules are close and visuals of dissimilar molecules are distant. The output was a probability that the two visuals showed the same molecule.

For training, we used back propagation with a history window and multiple back propagation passes, so as to emulate the fact that humans remember past consecutive problems and that humans update their internal models by reviewing the current problem along with the latest problem several times. Then, the effectiveness of the sequence is evaluated based on how well the cognitive model performs on a perceptual-fluency test, which is composed of a sample of perceptual-fluency problems drawn from a separate distribution of problems (i.e., training and test sequences contain different molecules). We use separate test and training distributions to ensure that we optimize the training sequence for learning of mappings among the visual features of the representations, rather than for memorization of translations for specific molecules.

In our prior work [38], we used data from novice undergraduate students in a chemistry course to develop the cognitive model. We then used a modified hill climb search algorithm to find an appropriate training sequence for thatmodel. Next, we compared this sequence to an expert-generated sequence in an experiment with participants from MTurk. Results showed that the machine-learned sequence yielded significantly higher gains in perceptual fluency than the expert-generated sequence.

3 Research Questions

While our prior work showed promising findings for MTurk participants, it remains an open question whether these benefits generalize to low-performing chemistry students. The MTurk participants in our prior study matched our target population because they had low prior knowledge about chemistry and little or no experience with the visual representations. However, their main motivation for participating in the study was to earn money, rather than to learn chemistry. Further, the MTurk study was not conducted in an educational setting. Hence, we address the following open questions:

Research question 1: Does the machine-learned sequence yield higher gains in perceptual fluency than an expert-generated sequence for chemistry students when embedded in instructional materials used in an undergraduate course?

Research question 2: Does the effectiveness of the machine-learned sequence depend on students' prior knowledge?

4 Methods

We address these questions in an experiment that compared the machine-learned sequence to an expert-generated sequence with chemistry undergraduate students with varying levels of prior knowledge enrolled in a chemistry course.

4.1 Participants

We conducted the experiment in a 300-level introductory chemistry course for undergraduates. While the course is open to freshmen and has a prerequisite of students having completed at least one 100-level chemistry course, many students enroll as seniors and have not taken chemistry since their freshman year. Hence, students have highly variable prior knowledge levels. Students received the perceptual-fluency training as a homework assignment with an intelligent tutoring system (ITS) (see Sect. 4.2). Forty students completed the assignment. Two students were excluded because they were statistical outliers on a pretest or posttest (see Sect. 4.4), yielding $N = 38$ students.

4.2 Chem Tutor: An ITS for Undergraduate Chemistry

The chemistry course used the Chem Tutor system for homework. Chem Tutor is an ITS that provides complex problems with individualized step-by-step guidance [4, 39]. Chem Tutor provides interactive instruction that introduces students to how visuals show chemistry concepts (see Fig. 3). In the assignment we used for this experiment, students received four instructional activities prior to the perceptual-fluency training.

The perceptual-fluency training of the assignment was structured as follows. First, students watched a 3-min video explaining that they would receive a large number of single-step problems in a row. The video explained that these problems served to train their perceptual fluency in quickly translating among visuals. Students were instructed not to overthink their answer but to intuitively decide if the two visuals showed the same molecule or not. Further, they were instructed that they would first receive a sequence of problems without feedback (i.e., pretest), then a sequence of problems with feedback (i.e., training), and finally problems without feedback (i.e., posttest).

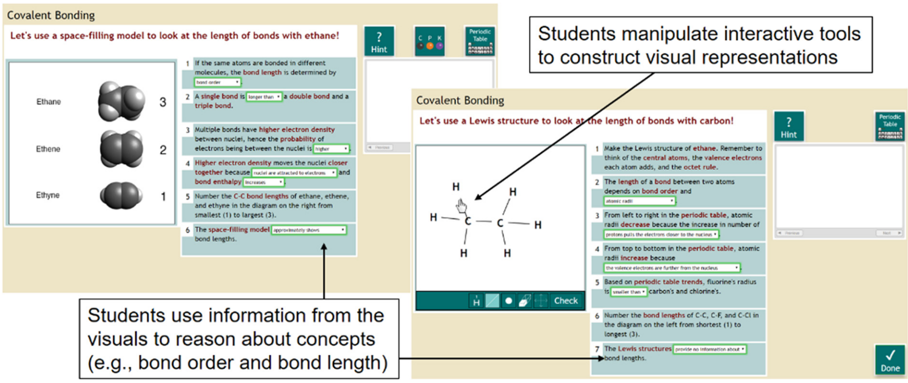

Fig. 3. Example sense-making activity in Chem Tutor.

4.3 Experimental Design

Students were randomly assigned to an expert-generated sequence or a machine-learned sequence that we used in the MTurk study [38]. Each sequence had sixty problems. To control for potential response biases that could affect learning, the number of problems that showed visuals of the same molecules was the same for both sequences.

The *expert-generated sequence* was created by a researcher with a decade of experience with perceptual-fluency trainings, using the principles that have been established by prior research on perceptual-fluency trainings we reviewed above. Specifically, problems were sequenced so they would draw attention to relevant visual features. To this end, consecutive problems often repeated one visual while changing the second visual. For example, if one problem presented visuals that showed different molecules (e.g., in the left of Fig. 2, the Lewis structure has more carbon atoms and the wrong bond order), the next problem might present visuals that showed the same molecules (e.g., the right of Fig. 2). To create such sequences, we randomly set the length of the subsequence that retained one visual to be 1–4 problems (i.e., either the first, second, third, or fourth problem would present visuals showing the same molecule). Then, we systematically varied visual features that play a role in chemistry learning, as determined by our prior research with novice students and chemistry experts [4, 40].

The *machine-learned sequence* was constructed using the machine-teaching approach described above. A qualitative inspection of the sequence reveals several differences to the expert-generated sequence that are worth highlighting. First, it does not repeat visuals across consecutive problems. Second, it contains problems that can be solved purely based on knowing which atoms the letters and colors in Lewis structures and space-filling models stand for; which is not one of the visual features the expert-generated sequence systematically varied. Third, it contains problems that can be solved by simply counting the number of atoms in the visuals; which is also not a visual feature that the expert-generated sequence aimed to draw attention to.

4.4 Measures

To assess students' gains in perceptual fluency, we used the same pretest and posttest as in our prior MTurk study [38]. As mentioned, pretest, training, and posttest problems were drawn from separate distributions to ensure that we assess learning of mappings among the visual features of the representations and not of memorization of translations between visuals of specific molecules. For brevity, the pretest contained only 20 problems; the posttest contained 40 problems. Students received no feedback on the test problems. Because perceptual fluency describes students' efficiency in seeing meaningful information in visuals, we computed efficiency scores for each test. Following prior work on efficiency measures [41], we computed perceptual-fluency scores as:

$$\text{perceptual - fluency score} = \frac{Z(\text{average correct responses}) - Z(\text{average time per problem})}{\sqrt{2}} \quad (1)$$

Further, to test if the effect of sequence depends on students' prior knowledge, we used the logs from the four interactive instruction activities that students completed prior to the perceptual-fluency problems. We computed prior-knowledge scores as the number of steps students answered correctly on the first attempt. Because the instruction activities ask students to answer questions about chemistry concepts based on the visuals, this measure assesses students' knowledge about how the visuals show concepts. We treated prior knowledge as a continuous variable in all analyses.

4.5 Procedure

Students were assigned to the Chem Tutor activities as homework in the second week of the semester, including the interactive instruction, pretest, perceptual-fluency training, and posttest. Students were given seven days to complete the assignment online.

5 Results

In the following analyses, we report p. η^2 effect sizes. Following Cohen [42], we consider p. η^2 of .01 to be a small effect, .06 a medium, and .14 a large effect.

5.1 Prior Checks

First, we checked for learning gains using repeated measures ANOVAs with pretest and posttest as dependent measures. Results showed large significant gains in perceptual fluency from the pretest to the posttest, $F(1,36) = 8.762$, $p = .005$, p. $\eta^2 = .196$.

Second, a multivariate ANOVA showed no significant differences between conditions on the perceptual-fluency pretest or prior knowledge on Chem Tutor's interactive instruction activities (Fs < 1). Further, there were no differences between conditions in terms of how much time students spent on the perceptual-fluency training ($F < 1$).

5.2 Effects of Sequence

To test if the machine-learned sequence yielded higher gains in perceptual fluency than an expert-generated sequence (research question 1), we used an ANCOVA model with condition as independent factor, perceptual-fluency pretest and prior knowledge as covariates, and perceptual-fluency posttest as dependent measure. To test if the effects depend on students' prior knowledge (research question 2), we added an interaction between condition and prior knowledge to the model. In line with prior research on aptitude-treatment interactions [43], we did not dichotomize prior knowledge but modeled the interaction between condition and the continuous prior-knowledge variable.

Results showed a medium-sized significant main effect of condition, $F(1,33) = 4.699$, $p = .037$, p. $\eta^2 = .125$, such that the machine-learned sequence yielded higher gains in perceptual fluency than the expert-generated sequence. The main effect was qualified by a medium-sized significant interaction of condition with prior knowledge, $F(1, 33) = 4.788$, $p = .036$, p. $\eta^2 = .127$. As shown in Fig. 4, the machine-generated sequence was more effective for students with lower prior knowledge, but the expert-generated sequence was more effective for students with higher prior knowledge.

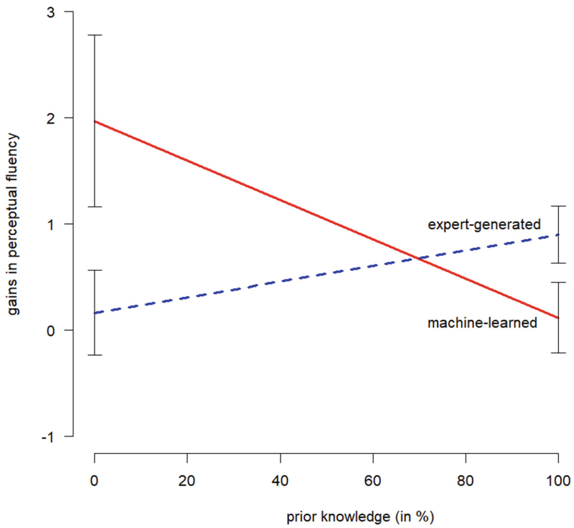

Fig. 4. Effect of machine-learned (red-solid) vs expert-generated (blue-dashed) sequence. The y-axis shows pre-post gains in perceptual-fluency scores based on the efficiency measure in Eq. (1). The x-axis shows prior knowledge. Error bars show standard errors of the mean. (Color figure online)

6 Discussion

Our goal was to investigate if a machine-learning approach enhances the effectiveness of perceptual-fluency trainings for low-performing students in a realistic educational context. We drew on our prior work on a machine-learning approach that was not subject to a potential blind spot bias due to experts' perceptual fluency in seeing meaning in visuals. Instead, it used a bottom-up approach to machine-learn a sequence of visuals based on data from novice chemistry students. Our prior work had established the effectiveness of this sequence for MTurk participants, who were not representative of students in a realistic educational context. The present findings replicate this effect in an undergraduate chemistry course and make several novel contributions.

First, our results show that the machine-learned sequence yield higher gains in perceptual fluency than an expert-generated sequence for students with lower prior knowledge. This finding shows that our machine-learning approach is an effective method for developing perceptual-fluency trainings that are attuned to the needs of students whose needs may not be obvious to instructional designers.

Second, our experiment makes new contributions to the perceptual fluency literature. In contrast to prior research, our findings suggest that perceptual-fluency trainings can be effective for students with low prior knowledge, but that these students require different types of such trainings. Our qualitative comparison of the machine-learned and expert-generated sequences suggests that students with low prior knowledge may benefit from sequences that draw attention to visual features that may seem obvious to experts, such as the mapping between letters and colors. Given that students in our experiment likely had some exposure to the visuals in prior chemistry courses, we think it is unlikely that they did not know that these features are important. Rather, they may not have been efficient at perceiving these features. Further, the machine-learned sequence did not repeat visuals across consecutive problems, whereas the expert-generated sequence did. Such repetitions assume that students recall the visuals from previous problems, which is cognitively demanding. Hence, students with low prior knowledge may benefit more from sequences that reduce cognitive load.

Third, we found that the expert-generated sequence is more effective for students with high prior knowledge. This finding replicates prior research on the effectiveness of expert-generated sequences for advanced students. Anew contribution of our findings is that we found that students' performance on prior instructional activities with visuals predicts if they have the prerequisite knowledge to benefit from an expert-generated sequence or if they should receive a sequence that was machine-learned based on data from novice students to prevent expert blind spot biases.

Our findings should be interpreted in the context of the following limitations. First, we focused on a specific set of visuals in chemistry. While we believe that the role of perceptual fluency in chemistry is representative of other STEM domains that rely heavily on visuals, future research should test if our findings generalize to other domains. Second, we did not contrast the characteristics of machine-learned and expert-generated sequences that may account for our results. For example, we did not test if the repetition of visual representations across problems is effective for students with high vs. low prior knowledge. Yet, our findings provide first indications that these characteristics may affect the acquisition of perceptual fluency, which can be systematically tested in future research. Third, because our sample size was relatively small, it is possible that additional smaller effects remained undetected. Finally, we assessed gains of perceptual fluency but not learning of content knowledge. Hence, future research should test whether gains in perceptual fluency for low-performing students translates into an enhanced ability to use the visual representations to learn content knowledge.

In sum, our experiment shows that a bottom-up approach to learn a sequence of visuals for perceptual-fluency trainings can help overcome potential biases resulting from an expert blind spot on the part of instructional designers. Such sequences are particularly effective for students with low prior knowledge. Further, our research provides new directions for future research to systematically investigate which characteristics enhance the acquisition of perceptual fluency for students with low prior knowledge.

Acknowledgements. This research was funded by NSF IIS 1623605 and by NSF ITP 1545481. We thank Yuzi Yu, Blake Mason, John Moore, Rob Nowak, Purav Patel, and Tim Rogers.

References

1. Gilbert, J.K.: Visualization: an emergent field of practice and inquiry in science education. In: Gilbert, J.K., Reiner, M., Nakhleh, M.B. (eds.) Visualization: Theory and Practice in Science Education, vol. 3, pp. 3–24. Springer, Dordrecht (2008). https://doi.org/10.1007/978-1-4020-5267-5_1
2. Ainsworth, S.: The educational value of multiple-representations when learning complex scientific concepts. In: Gilbert, J.K., Reiner, M., Nakama, A. (eds.) Visualization: Theory and Practice in Science Education, pp. 191–208. Springer, Dordrecht (2008). https://doi.org/10.1007/978-1-4020-5267-5_9
3. NRC: Learning to Think Spatially. National Academies Press, Washington, D.C. (2006)
4. Rau, M.A., Michaelis, J.E., Fay, N.: Connection making between multiple graphical representations: a multi-methods approach for domain-specific grounding of an intelligent tutoring system for chemistry. Comput. Educ. **82**, 460–485 (2015)
5. Ainsworth, S.: DeFT: a conceptual framework for considering learning with multiple representations. Learn. Instr. **16**, 183–198 (2006)
6. McElhaney, K.W., Chang, H.Y., Chiu, J.L., Linn, M.C.: Evidence for effective uses of dynamic visualisations in science curriculum materials. Stud. Sci. Educ. **51**, 49–85 (2015)
7. Rau, M.A.: Conditions for the effectiveness of multiple visual representations in enhancing STEM learning. Educ. Psychol. Rev. **29**, 717–761 (2017)
8. Rau, M.A., Aleven, V., Rummel, N.: Successful learning with multiple graphical representations and self-explanation prompts. J. Educ. Psychol. **107**, 30–46 (2015)
9. Seufert, T., Brünken, R.: Cognitive load and the format of instructional aids for coherence formation. Appl. Cogn. Psychol. **20**, 321–331 (2006)
10. Bodemer, D., Faust, U.: External and mental referencing of multiple representations. Comput. Hum. Behav. **22**, 27–42 (2006)
11. Kellman, P.J., Massey, C.M.: Perceptual learning, cognition, and expertise. In: Ross, B.H. (ed.) The Psychology of Learning and Motivation, vol. 558, pp. 117–165. Elsevier Academic Press, New York (2013)

12. Goldstone, R.L., Barsalou, L.W.: Reuniting perception and conception. Cognition **65**, 231–262 (1998)
13. Kellman, P.J., Garrigan, P.B.: Perceptual learning and human expertise. Physics Life Rev. **6**, 53–84 (2009)
14. Goldstone, R.L., Schyns, P.G., Medin, D.L.: Learning to bridge between perception and cognition. Psychol. Learn. Motiv. **36**, 1–14 (1997)
15. Gibson, E.J.: Perceptual learning in development: some basic concepts. Ecol. Psychol. **12**, 295–302 (2000)
16. Goldstone, R.: Perceptual Learning. Academic Press, San Diego (1997)
17. Rau, M.A., Wu, S.P.W.: Support for sense-making processes and inductive processes in connection-making among multiple visual representations. Cogn. Instr. **36**, 361–395 (2018)
18. Rau, M.A.: Sequencing support for sense making and perceptual induction of connections among multiple visual representations. J. Educ. Psychol. **110**, 811–833 (2018)
19. Schooler, J.W., Fiore, S., Brandimonte, M.A.: At a loss from words: verbal overshadowing of perceptual memories. Psychol. Learn. Motiv. Adv. Res. Theor. **37**, 291–340 (1997)
20. Goldstone, R.L., Landy, D.H., Son, J.Y.: The education of perception. Top. Cogn. Sci. **2**, 265–284 (2010)
21. Gibson, E.J.: Principles of Perceptual Learning and Development. Prentice Hall, New York (1969)
22. Chase, W.G., Simon, H.A.: Perception in chess. Cogn. Psychol. **4**, 55–81 (1973)
23. Richman, H.B., Gobet, F., Staszewski, J.J., Simon, H.A.: Perceptual and memory processes in the acquisition of expert performance: the EPAM model. In: Ericsson, K.A. (ed.) The Road to Excellence? The Acquisition of Expert Performance in the Arts and Sciences, Sports and Games, pp. 167–187. Erlbaum Associates, Mahwah (1996)
24. Taber, K.S.: Revisiting the chemistry triplet: drawing upon the nature of chemical knowledge and the psychology of learning to inform chemistry education. Chem. Educ. Res. Pract. **14**, 156–168 (2013)
25. Koedinger, K.R., Corbett, A.T., Perfetti, C.: The knowledge-learning-instruction framework: bridging the science-practice chasm to enhance robust student learning. Cogn. Sci. **36**, 757–798 (2012)
26. Shanks, D.: Implicit learning. In: Lamberts, K., Goldstone, R. (eds.) Handbook of Cognition, pp. 202–220. Sage, London (2005)
27. Massey, C.M., Kellman, P.J., Roth, Z., Burke, T.: Perceptual learning and adaptive learning technology - developing new approaches to mathematics learning in the classroom. In: Stein, N.L., Raudenbush, S.W. (eds.) Developmental Cognitive Science Goes to School, pp. 235–249. Routledge, New York (2011)
28. Rau, M.A., Aleven, V., Rummel, N.: Interleaved practice in multi-dimensional learning tasks: which dimension should we interleave? Learn. Instr. **23**, 98–114 (2013)
29. Rau, M.A., Aleven, V., Rummel, N., Pardos, Z.: How should intelligent tutoring systems sequence multiple graphical representations of fractions? A multi-methods study. Int. J. Artif. Intell. Educ. **24**, 125–161 (2014)
30. Bradley, J.-C., Lancashire, R., Lang, A., Williams, A.: The spectral game: leveraging open data and crowdsourcing for education. J. Cheminform. **1**, 1–10 (2009)
31. Rau, M.A., Aleven, V., Rummel, N.: Supporting students in making sense of connections and in becoming perceptually fluent in making connections among multiple graphical representations. J. Educ. Psychol. **109**, 355–373 (2017)
32. Kellman, P.J., et al.: Perceptual learning and the technology of expertise: studies in fraction learning and algebra. Pragmat. Cogn. **16**, 356–405 (2008)

33. Wise, J.A., Kubose, T., Chang, N., Russell, A., Kellman, P.J.: Perceptual learning modules in mathematics and science instruction. In: Hoffman, P., Lemke, D. (eds.) Teaching and Learning in a Network World, pp. 169–176. IOS Press, Amsterdam (2000)
34. Kellman, P.J., Massey, C.M., Son, J.Y.: Perceptual learning modules in mathematics: enhancing students' pattern recognition, structure extraction, and fluency. Top. Cogn. Sci. **2**, 285–305 (2009)
35. Nathan, M.J., Koedinger, K.R., Alibali, M.W.: Expert blind spot: when content knowledge eclipses pedagogical content knowledge. In: Chen, L., et al. (eds.) Proceedings of the Third International Conference on Cognitive Science, pp. 644–648. USTC Press, Beijing (2001)
36. Chin, J.M., Schooler, J.W.: Why do words hurt? Content, process, and criterion shift accounts of verbal overshadowing. Eur. J. Cogn. Psychol. **20**, 396–413 (2008)
37. Zhu, X.: Machine teaching: an inverse problem to machine learning and an approach toward optimal education. In: AAAI Conference on Artificial Intelligence (2015)
38. Sen, A., et al.: Machine beats human at finding the optimal sequence of visual representations for students learning of perceptual fluency. In: Boyer, K.E., Yudelson, M. (eds.) Proceedings of the 11th International Conference on Educational Data Mining, pp. 137–146. International Educational Data Mining Society, Buffalo (2018)
39. VanLehn, K.: The relative effectiveness of human tutoring, intelligent tutoring systems and other tutoring systems. Educ. Psychol. **46**, 197–221 (2011)
40. Rau, M.A., Mason, B., Nowak, R.: How to model implicit knowledge? Use of metric learning to assess student perceptions of visual representations. In: Barnes, T., Chi, M., Feng, M. (eds.) Proceedings of the 9th International Conference on Educational Data Mining, pp. 199–206. International Educational Data Mining Society, Raleigh (2016)
41. Van Gog, T., Paas, F.: Instructional efficiency: revisiting the original construct in educational research. Educ. Psychol. **43**, 16–26 (2008)
42. Cohen, J.: Statistical Power Analysis for the Behavioral Sciences. Lawrence Erlbaum Associates, Hillsdale (1988)
43. Park, O., Lee, J.: Adaptive instructional systems. In: Jonassen, D.H. (ed.) Handbook of Research for Educational Communications and Technology, pp. 651–658. Erlbaum, Mahwah (2003)