Issues in Information Systems
Volume 22, Issue 2, pp. 41-51, 2021

DOI: https://doi.org/10.48009/2_iis_2021_41-51

Exploring the use of sub-goals to train students to solve novel problems in a SQL context

Brian Kimball Dunn, *Utah State University, brian.dunn@usu.edu*
David Olsen, *Dixie State University, david.olsen@dixie.edu*
Reagan Rice, *Utah State University, reaganrice2323@gmail.com*

Abstract

Within this study, we address the question of how the use of sub-goals within a technical advanced database course promotes students' abilities to address novel problems that rely on their technical skills. We report on a learning module developed with this in mind. The module requires students to apply known skills, such as CAST and CROSS JOIN functions and SQL variable declarations as well as newly presented content regarding SQL Server's geography-related functions and data types, toward solving a novel problem. The assignment first introduced students to ways of using knowledge of SQL to address three specified sub-goals within a geography-related database problem. After practicing the steps required to address these sub-goals, students were then asked to complete a new geography-related database problem that did not closely match the content or context of the practice example.

Keywords: SQL, Teaching Novel Problem-Solving, Geography, Sub-Goals

Introduction

Organizations expect their employees to deal effectively with novel problems on a daily basis. This is particularly true for employees in information technology-related positions, where rapid development and application of new technologies are expected as is the drive to attain better value from current technology investments. Therefore, students who learn how to apply skills to problems that are novel to them should be better prepared to help their organizations and advance their careers once they are in the workforce.

Successfully teaching students to solve novel problems, however, presents a difficult challenge. A considerable body of research shows that students struggle with problems that do not closely mimic the context and content of examples they have already learned (e.g., Pirolli & Anderson, 1985; Reed et al., 1990; Ross, 1989; Wang & Adesope, 2017). Novel problems present students with a large search space to consider when seeking an appropriate approach to solving the problem and are ill-structured (Chappell et al., 2013; Newell & Simon, 1972; Simon & Reed, 1976). Past studies show that students learn more effectively how to solve novel problems when they are taught to work toward specific sub-goals (Atkinson et al., 2000; Catrambone, 1996; Morrison et al., 2016). Sub-goals — intermediate steps required when solving a problem — provide structure and enable learners to narrow the search space required to address the problem (Catrambone, 1994).

This study addresses the question of how the use of sub-goals within a technical advanced database course promotes students' abilities to address novel problems that rely on their technical skills. We report on a learning module developed with this in mind. The module requires students to apply known skills, such as CAST and CROSS JOIN functions and SQL variable declarations as well as newly presented content regarding SQL Server's geography-related functions and data types, toward solving a novel problem. The assignment first introduced students to ways of using knowledge of SQL to address three specified sub-goals within a geography-related database problem. After practicing the steps required to address these sub-goals, students complete a new geography-related database problem that did not closely match the content or context of the practice example.

Using this approach, many students reported that they felt enabled to solve the novel problem. Indeed, based on the work submitted by students, they demonstrated having applied their skills effectively to the novel situation provided. This suggests that this sub-goal approach succeeded in narrowing the search space required for students to address the problem and enabled them to better address the challenge. The results of this study indicate that a clear focus on sub-goals, including learning of each sub-goal's purpose and fit within the greater problem, represents a useful approach when teaching the use of technology within a problem-solving context.

In the next section, we discuss novel problems and the theory underlying sub-goals and problem spaces. Then, we describe the module that was taught and how it utilized these principles to facilitate learning. We then discuss our findings based on students' experiences with the module and the implications of these findings for classroom use and further research.

Background

Increasingly, career success depends an employee's ability to harness complex skillsets to critically resolve novel problems (Altstadt, 2010). In this study, we consider a *novel problem* to be any problem or situation that differs from other problems and situations faced by an individual (e.g., a student) (Chan & Carlson, 2017). This contrasts with the concept of an *isomorphic problem,* which is similar in form and context to what has already been performed (Chi et al., 1989). This is particularly true for IT professionals, given the complexity involved in finding solutions to many information technology-related problems. Such problems involve myriad systems, languages, protocols, etc. that must be understood and brought together to accomplish the objective.

However, people are not naturally adept at addressing novel problems. Considered within the context of a learning environment, individuals struggle when asked to address a problem not presented within a familiar frame. They find it difficult to generalize solutions out from the specific context and examples used to teach the given skills (Chi et al., 1989; Reed et al., 1990) and to develop a structure to help themselves address the problem (Chappell et al., 2013). Most information systems faculty have experienced this when teaching complex concepts. A student may understand how to write a SELECT query to output customer names, ages, and phone numbers from a customer table, but then struggle to write a very similar query to output product names, categories, and prices from a product-related table. Similarly, beginning programming students may be able to use a for-loop to print out elements in a list, but then not be able to apply that same knowledge to print out lines from a text object.

Given this, how can we better teach students to solve novel problems, both in the classroom and once in the workforce? More specifically, how can we teach complex information technology skills, such as SQL querying, to facilitate its use in novel situations?

Problem Spaces

To investigate these questions, we refer to the idea of *problem spaces* (Newell & Simon, 1972), the mental representation of a problem, including the initial state, the goal state, the set of actions and tools used to progress toward the eventual goal, and any constraints placed on the pathway to the solution (Erickson & Jones, 1978; Newell, 1980; Newell & Simon, 1972). When given a problem, the problem-solver searches within the problem space looking for steps that can help advance the problem toward its solution. In this process, the solver evaluates possible actions, then, based on the evaluation, undertakes those perceived to advance the problem toward its resolution. After this, the problem-solver appraises the problem's new state and considers additional actions (Newell, 1980).

These novel problems present a large, ill-structured problem space (Chappell et al., 2013). Given a novel problem with which the problem-solver is relatively unfamiliar, the problem space contains many more possible alternatives and intermediate states to consider. For instance, consider tasking someone who has no experience with carpentry to undertake a do-it-yourself kitchen remodel. This person would need to search an almost infinite problem space to come up with what might be a good first step. During this process and given that the problem-solver is altogether new to the problem, it is likely that the first step identified will be an incorrect or sub-optimal one.

Compare this to a more familiar problem space, where the problem-solver has high confidence in knowing that the solution requires a combination of a much smaller set of possible steps. Think of the problem of an experienced motorist approaching a stoplight. This presents a relatively smaller problem space — given the condition of the stoplight, the motorist only needs to search a limited set of possible actions (e.g., slow down, speed up, or continue at the same rate).

In this comparison, the first problem is novel and complex and will take considerably longer to resolve (even ignoring the time and effort to complete the physical carpentry work) given that the problem-solver must establish their own structure for the problem before being able to pursue its solution (McGraw & Patterson, 2019). The second, given its clearer structure and much more limited problem space will be considerably easier.

In moving from the initial problem state toward the solution to the problem, a problem-solver's actions change the current state (Newell, 1980). Each of these state changes may be a necessary part of the solution *or* may simply be a state that, while possible, is not needed (Newell, 1980). In a larger problem space, then, it is more likely that the problem-solver will undertake actions that arrive at a possible, but not a necessary, problem state. It is difficult for the problem-solver to identify the correct next step, therefore the problem-solver will make mistakes (Simon, 1978). This occurs, in part, due to the cognitive load experienced in the course of addressing the problem (Sweller, 1988; Wang & Adesope, 2017).

Self-Explanation and Sub-Goals

Within the pedagogical context, past work demonstrates that those students who generate many explanations to refine and expand those actions performed in completing example problems and relate them to other principles perform better at applying their skills to novel problems (Chi et al., 1989). A key part of this process is called *self-explanation*, the process by which an individual recognizes and thereby gains an understanding of the fundamental characteristics and relationships of a principle (Rittle-Johnson et al., 2017; Simon, 1979). Through self-explanation, learners develop an example-independent knowledge, which can then be applied to problems in other contexts — that is, novel problems.

Further, the fundamental understanding gained through self-explanation allows the learner to narrow the problem space, thus facilitating better and faster solutions to novel problems with lower cognitive load (Wang & Adesope, 2017). However, while some learners seek self-explanation natively when new ideas or skills presented, this is not the case for all. Thus, an approach that encourages self-explanation can improve the novel problem-solving abilities for all learners.

Self-explanation can be facilitated through the use of *sub-goals* (Catrambone, 1994; Huang et al., 2017; Newell, 1980; Richland et al., 2017) — milestones or smaller achievements that contribute to the overall problem-solving process. For instance, in the kitchen remodel example above, sub-goals may include steps such as drawing up plans, consulting local building codes, completing demolition of the existing kitchen, ordering cabinets, etc. These sub-goals act as guides and provide structure to the problem, signposting the steps that the problem-solver should take and thereby reducing unnecessary cognitive load and leading to better solutions and higher motivation (Renkl & Atkinson, 2003; Richland et al., 2017).

For a number of reasons, however, learners struggle to identify appropriate sub-goals when left to their own devices (Catrambone, 1994; Singley & Anderson, 1989), even after an instructor demonstrates solutions to a given problem. When initially working through examples and practice problems, extraneous details distract learners who then fail to notice, let alone focus on, the fundamental concepts included in the example (Morrison et al., 2016; Sweller, 2010; Wang & Adesope, 2017). A linear regression analysis problem presented within the context of predicting the number of defects that occur in a jack-in-the-box factory, for instance, can result in students remembering the jack-in-the-boxes, but not how linear regression was used to find a solution.

However, by demonstrating the solution to problems while using and emphasizing the presence and function of sub-goals, students learn to self-explain the function of the sub-goals and thereby narrow the problem space. This allows them to perform better when applying skills to novel problems (Catrambone, 1994; Margulieux et al., 2012; Morrison et al., 2016; Wang & Adesope, 2017).

Teaching Sub-Goals Effectively

As noted, when learning a new problem-solving procedure, students tend to memorize the superficial aspects of the procedure rather than obtain deeper conceptual knowledge through self-explanation (Catrambone, 1998). While teaching such rote procedures has been shown to be effective during the early stages of skill development (Atkinson et al., 2000), it does not by itself facilitate learner ability to solve novel problems. Therefore, when teaching complex procedures, calling out the presence and function of sub-goals within the procedure causes students to focus on the procedure's structural features (Atkinson et al., 2000), thereby reducing the cognitive load required to solve the problem (Renkl & Atkinson, 2003).

Further, when explaining sub-goals to learners, learners must understand the conditions and consequences of each within the context of the overall procedure (Chi et al., 1989). This enables the learner to know the conditions when the sub-goal applies (or doesn't apply), the consequences of completing the sub-goal, and the relationship between the sub-goal and the overall procedure (Chi et al., 1989). Explicitly labeling and calling attention to these sub-goals is more effective than relying on their latent presence (Catrambone, 1998; Morrison et al., 2016).

Illustration

To evaluate the usefulness of this sub-goal approach within the context of a complex IT-related skill, we created and tested the use of a new SQL module taught in a senior-level advanced database course. In this module, students learned the use of geographic features within SQL Server and were expected to apply these features, combined with existing knowledge of SQL, to solve a novel problem.

This context was conducive to this approach given that (a) at this point in their learning, they have acquired a number of skills, (b) the course itself exposes them to new, more-esoteric concepts, and (c) course content is readily applicable to a myriad of different problem types. Indeed, this specific geography module requires students to learn to use geographic data types that have become integral to numerous real-world software applications such as GIS, Google Maps, Uber, and a number of other practical applications.

Prior Learning

This module was part of an advanced database course, for which the introductory database course was a prerequisite. In that prerequisite course, students learned how to write SQL queries to interact with and manipulate data through Microsoft SQL Server Management Studio. In the advanced database course and prior to the introduction of this module, students mastered skills in many areas, including subqueries, recursive joins, common table expressions, window functions, and PL/SQL. Particularly relevant to this module, students also learned about SQL variable declarations, recursive joins, table updating, and the SQL Server CAST function. Prior to the beginning of the module, we also assumed that students had some basic knowledge of geography, particularly the use of latitude and longitude for describing locations.

Module Description

Throughout the course, modules were taught through a consistent format. In this format, students would first be instructed on new content, then work through an exercise along with the instructor that integrated the new content along with prior learned content, then complete a new assignment on their own. The geography module used in our study followed this same format.

The in-class portion of the module consisted of two stages, the example demonstration (Stage 1) and the novel application (Stage 2). For this module, we identified three sub-goals that were emphasized in Stage 1 and that could then be applied during Stage 2: (1) locations, (2) declarations, and (3) calculations. Sub-goal 1 (locations) requires specifying as coordinates (longitude and latitude) the locations required to solve the problem. Sub-goal 2 (declarations) requires loading (declaring) variables into SQL using appropriate SQL methods so as to relate location names to the given coordinates. Finally, sub-goal 3 (calculations) requires using commands to calculate the distances requested in the problem (initially in meters, then converted to miles).

Stage 1: Example

In the example stage, students sub-divided into groups. They were initially asked to complete sub-goal 1 by identifying the coordinates of several well-known locations. Specifically, the groups identified the latitude and longitude coordinates for their classroom building as well as six other locations: the Palace of Versailles, the Taj Mahal, Winter Palace in St. Petersburg, Cristo Redentor in Rio de Janeiro, Surfer's Paradise in Australia, and a sixth location of their choosing. While addressing the requirements of sub-goal 1, it also provided a chance for students to reacquaint themselves with the ideas of latitude and longitude in general, something that many students had not thought about in several years.

Once the students had attained the coordinate pair for each location, with help from the instructor and TA where needed, they then completed sub-goal 2 by using the appropriate SQL methods to declare each pair and assign it a corresponding location name. Students were already knowledgeable on the syntax for variable declaration; in this step, they combined this with the newly presented syntax specific to geographic data types. This resulted in the generalized code shown below.

```
DECLARE @locationName geography;
SET @locationName = geography::Point(Lat, Long, 4326);
```

The first line of code declares a variable (e.g., @versailles, @tajmahal) of data type geography. This data type, specific to SQL Server, allows the specification of both latitude and longitude to be applied to the variable declared. Then, the second line sets the value of the variable to equal the point specified by the given values for "Lat" and "Long", with latitude measured in decimal degrees north latitude and longitude in decimal degrees east longitude (locations in the western hemisphere, then, had negative longitude values). For example, to declare the variable for the Winter Palace location, students used a query similar to the following.

```
DECLARE @winterpalace geography;
SET @winterpalace = geography::Point(59.9404, 30.3138, 4326);
```

Note that the third parameter in the above example, shown with a value of 4326, denotes the spatial reference identifier (SRID) of the geography instance to be returned. SRID 4326 refers to the World Geodetic System 84 (WGS84), which serves as the standard for geographic mapping.

Once all locations' coordinates were assigned to corresponding variables (i.e., once sub-goal 2 was completed), sub-goal 3 required students to then calculate the distances from their classroom on campus to each of the six other locations identified. For this, they applied the following approach for calculating and outputting distances between two variables of the geography data type and outputting the distance in miles.

```
SELECT @location1.STDistance(@location2)/1609.344 AS [Distance from Loc1 to
Loc2];
```

In this query, students learned to use the STDistance operator with the syntax given above. However, the STDistance operator measures distance in meter units; the problem required an answer in miles, thus requiring students to divide the result by 1609.344 to complete the conversion. For a query to output the distance from the classroom building (Huntsman Hall) to the Winter Palace in miles, then, students needed to write a query similar to the following.

```
SELECT @winterpalace.STDistance(@huntsmanhall)/1609.344 AS [Distance from
Huntsman Hall to Winter Palace];
```

Stage 2: Novel Application

After completing this example stage (Stage 1), students moved to the novel application stage (Stage 2). In this second stage, students applied the three sub-goals from Stage 1 to solve a novel problem requiring them to apply what they had learned about geographic functions within SQL Server. Students used their skills to calculate and report the distances between all named cities and communities within the state of Utah, a total of 364 locations. Not only did this stage require students to apply the module-specific skills using geography methods and functions, but it also required them to pair these with their pre-existing knowledge of table updating, recursive joins, and the CAST function. Thus, Stage 2 represented a truly novel context in which to apply their new knowledge using the specified sub-goals.

For sub-goal 1 (location), students familiarized themselves with the longitude and latitude coordinates of all cities located in the state of Utah. Once completed, they moved to sub-goal 2 (declaration) by loading these location names and coordinates into a new table. To further familiarize themselves with the locations in the dataset, we encouraged them to write a query that displayed the latitude, longitude, geolocation, city name, and state for each city. This also required students to rely on their knowledge of the CAST command to ensure that the geolocation field displayed correctly.

```
SELECT Long, Lat,
       CAST(GeoLocation.ToString()) AS CHAR(25)) AS [Location],
       City, StateName
FROM UtahCities;
```

The query yielded output such as the following (only a part of the output is shown).

Long	Lat	Location	City	StateName
-112.7180000000000	39.3966000000000	POINT (-112.718 39.3966)	Abraham	Utah
-112.7938000000000	38.2583000000000	POINT (-112.7938 38.2583)	Adamsville	Utah
-111.7728000000000	40.4621000000000	POINT (-111.7728 40.4621)	Alpine	Utah
-111.6227000000000	40.5817000000000	POINT (-111.6227 40.5817)	Alta	Utah
-110.2886000000000	40.3586000000000	POINT (-110.2886 40.3586)	Altamont	Utah

Figure 1: Output of Latitude and Longitude Query

Then, to complete the problem, students addressed sub-goal 3 (calculation) by applying their knowledge to create a table that showed each city in the dataset and its calculated distance in miles to every other city in the dataset. For this, students applied not only the skills demonstrated in the stage 1 example stage (i.e., calculating distance in miles between two defined points), but also needed to integrate this with many other concepts learned earlier in the semester. A query to accomplish this final step looked like this:

```
SELECT c1.city, c2.city,
       c1.GeoLocation.STDistance(c2.GeoLocation)/1609.344  AS  [Distance  in
Miles]
FROM UtahCities AS c1 CROSS JOIN UtahCities AS c2
WHERE c1.City <> c2.City;
```

This query's output looked similar to the following (only a part of the output is shown).

city	city	Distance in Miles
Abraham	Adamsville	78.6261248744002
Abraham	Alpine	89.0186054746141
Abraham	Alta	100.318286748925
Abraham	Altamont	145.186641165331
Abraham	Alton	136.111392042617
Abraham	Amalga	175.130277622662
Abraham	American Fork	83.8526996344523
Abraham	Aneth	245.321850938616
Abraham	Angle	88.6830007818394
Abraham	Annabella	59.3953540182456

Figure 2: Output of Distance Calculation Query

Discussion

This study describes a successful application of the sub-goal approach to helping students solve novel problems within a technical IT context. In it, we explored the use of sub-goals to help facilitate students' successfully solving a SQL-related problem in a geography context. In the first stage of the module, students solved a distance-finding problem by addressing three different sub-goals, which corresponded with the ideas of location (sub-goal 1), declaration (sub-goal 2), and calculation (sub-goal 3). Aware of these sub-goals as a way to address such problems, students then applied the same sub-goals to a novel distance-finding problem that included some substantial differences.

The novel problem given presented students with a very large problem space. Without knowledge of the specified sub-goals, students may have been unsure as to how to proceed and which of their learned SQL skills they should apply. This would have led to the problem evolving into new states without progressing toward a satisfying conclusion. By applying the sub-goals, however, students narrowed the problem space, limiting their search for appropriate actions to those specified by the sub-goals. Further, these sub-goals helped students toward acquiring expertise by providing them narrowed objects that they could relate to each other, thus facilitating the process of self-explanation.

The context used in this demonstration of the sub-goal approach reflects a typical context that technically-oriented students will face once in the working world. Advanced SQL topics such as this one are similar to other technical skills, such as programming, data analysis, digital forensics, etc., where there are number of interrelated skills needed and the relationships between these skills must be understood to solve problems successfully. The demonstration of this sub-goal approach within this SQL context, then, suggests that similar approaches would work in similarly complex technical domains. Following similar procedures will enable students to apply their learned skills not only in the classroom setting and in the canned contexts in which they sometimes taught, but to also understand how to apply them when faced with new problems that their skills should help them resolve.

Recommendations

With the classroom application of this sub-goal approach, we followed canonical best practices (e.g., Catrambone, 1998; Morrison et al., 2016). Given the success of this, we recommend that others using sub-goals to teach similar complex skill applications do so as well. In particular, in applying this technique, we recommend (1) making explicit to students the presence and function of the sub-goal approach, (2) explicitly labeling and calling attention to each of the sub-goals, and (3) clearly explaining the conditions and consequences of each sub-goal within the context of the overall problem-solving process. We highlighted this three-step approach to students. In doing so, we indicated that each of these steps was useful in addressing geography-related SQL problems and that the presence of the steps would, when understood and learned, help in applying their skills to novel problems in the future.

Next, each of the sub-goals was given a clear label. To this end, we used the terms “location”, “declaration”, and “calculation” to describe each of the three sub-goals respectively. In adding these labels to represent the specified sub-goals, the sub-goals themselves became easier for students to recollect and apply to novel situations.

Finally, we discussed the conditions and consequences of each of the sub-goals within the context of the geography problem-solving context, thus facilitating the process of self-explanation (Chi et al., 1989). These conditions included understanding the state(s) when a given sub-goal should be pursued as well as the state(s) that should result from the pursuit of each sub-goal.

Future Research

The success of this approach suggests the need for further studies. For one, the premise of the sub-goal approach relies on its narrowing the problem space, thus making application of known skills to new contexts easier. Implicit here is that a narrowed problem space decreases the cognitive load and the number of mistakes made (Sweller, 2010; Wang & Adesope, 2017). However, in this example, we did not explicitly measure the degree of cognitive load experienced by students when using the sub-goal approach. Further, while all students came up with credible solutions to the Stage 2 problem, we did not measure the extent to which students took "wrong turns" on the way to coming up with their solutions. Measuring these would help us better understand the mechanism by which sub-goals facilitate novel problem solving within this context.

Further, a two-case comparative study would give us a stronger sense of how much this approach actually helps students complete the novel problem. While our experience as educators observing students and comments from students themselves suggest the success of the approach, a two-case comparison, wherein a control group is not given explicit sub-goals during Stage 1, would provide us information about the degree to which this approach helps apply learning to novel situations. Further, combining this with measurements of cognitive load and the number of mistakes made (i.e., task accuracy) would give more insight as to the mechanism by which sub-goals facilitate problem-solving.

Lastly, the sub-goals used in this example were based on what appeared to be clear steps in this specific problem-solving process. However, we did not examine whether more- or less-complex sub-goals yielded better performance. A further investigation experimenting with the degree of complexity inherent within sub-goals could yield new insights into how to best apply this approach.

Conclusion

Solving novel problems presents a significant challenge. As such, identifying credible approaches to guide students to the acquisition of skills in this area becomes an important task for technical skill educators. To that end, we examined the processes involved in helping students learn to solve hard problems in new contexts using the lens of sub-goals. From this, we saw that sub-goals were useful in solving these novel problems in that they enabled students to focus on a clear structure and a narrowed problem space.

This sub-goal paradigm was applied to an SQL geography problem because it is both novel and pervasive in analytics circles. Anecdotal evidence suggests students responded positively to the entire exercise of geographical data and the various tasks they were asked to perform. This initial study portends an exciting set of future research studies that can rigorously examine the advantages of the sub-goal approach.

References

Altstadt, D. (2010). Building opportunity: How states can leverage capital and infrastructure investments to put working families on a path to good jobs. *The Working Poor Families Project*.

Atkinson, R. K., Derry, S. J., Renkl, A., & Wortham, D. (2000). Learning from examples: Instructional principles from the worked examples research. *Review of Educational Research*, *70*(2), 181–214.

Catrambone, R. (1994). Improving examples to improve transfer to novel problems. *Memory & Cognition*, *22*(5), 606–615.

Catrambone, R. (1996). Generalizing solution procedures learned from examples. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(4), 1020.

Catrambone, R. (1998). The subgoal learning model: Creating better examples so that students can solve novel problems. *Journal of Experimental Psychology: General*, *127*(4), 355.

Chan, T., & Carlson, M. C. (2017). Building upon the CALLA framework: The intersection of novel problem solving and applied wisdom. *Human Development*, *59*(6), 366.

Chappell, J., Cutting, N., Apperly, I. A., & Beck, S. R. (2013). The development of tool manufacture in humans: What helps young children make innovative tools? *Philosophical Transactions of the Royal Society B: Biological Sciences*, *368*(1630), 1–10.

Chi, M. T., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, *13*(2), 145–182.

Erickson, J. R., & Jones, M. R. (1978). Thinking. *Annual Review of Psychology*, *29*(1), 61–90.

Huang, S., Jin, L., & Zhang, Y. (2017). Step by step: Sub-goals as a source of motivation. *Organizational Behavior and Human Decision Processes*, *141*, 1–15.

Margulieux, L. E., Guzdial, M., & Catrambone, R. (2012). *Subgoal-labeled instructional material improves performance and transfer in learning to develop mobile applications*. 71–78.

McGraw, R., & Patterson, C. L. (2019). Establishing Problem Spaces and Boundaries during Small Group Work on a Contextualized Problem. *Journal of Mathematics Education*, *12*(2), 1–17.

Morrison, B. B., Margulieux, L. E., Ericson, B., & Guzdial, M. (2016). *Subgoals help students solve Parsons problems*. 42–47.

Newell, A. (1980). Reasoning, Problem Solving and Decision Processes: The Problem Space as a Fundamental Category. In R. S. Nickerson (Ed.), *Attention and Performance: Vol. VIII*. Erlbaum.

Newell, A., & Simon, H. A. (1972). *Human problem solving* (Vol. 104). Prentice-Hall Englewood Cliffs, NJ.

Pirolli, P. L., & Anderson, J. R. (1985). The role of learning from examples in the acquisition of recursive programming skills. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, *39*(2), 240.

Reed, S. K., Ackinclose, C. C., & Voss, A. A. (1990). Selecting analogous problems: Similarity versus inclusiveness. *Memory & Cognition*, *18*(1), 83–98.

Renkl, A., & Atkinson, R. K. (2003). Structuring the transition from example study to problem solving in cognitive skill acquisition: A cognitive load perspective. *Educational Psychologist*, *38*(1), 15–22.

Richland, L. E., Begolli, K. N., Simms, N., Frausel, R. R., & Lyons, E. A. (2017). Supporting mathematical discussions: The roles of comparison and cognitive load. *Educational Psychology Review*, *29*(1), 41–53.

Rittle-Johnson, B., Loehr, A. M., & Durkin, K. (2017). Promoting self-explanation to improve mathematics learning: A meta-analysis and instructional design principles. *ZDM*, *49*(4), 599–611.

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(3), 456.

Simon, H. A. (1978). Rationality as process and as product of thought. *The American Economic Review*, *68*(2), 1–16.

Simon, H. A. (1979). Problem solving and education. In D. Tuma & F. Reif (Eds.), *Problem Solving and Education: Issues in Teaching and Research*. Erlbaum.

Simon, H. A., & Reed, S. K. (1976). Modeling strategy shifts in a problem-solving task. *Cognitive Psychology*, *8*(1), 86–97.

Singley, M. K., & Anderson, J. R. (1989). *The transfer of cognitive skill*. Harvard University Press.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, *12*(2), 257–285.

Sweller, J. (2010). Element interactivity and intrinsic, extraneous, and germane cognitive load. *Educational Psychology Review*, *22*(2), 123–138.

Wang, Z., & Adesope, O. (2017). Do focused self-explanation prompts overcome seductive details? A multimedia study. *Journal of Educational Technology & Society*, *20*(4), 47–57.