Effects of Goal Specificity on a Search in a Hypothesis Space and an Instance Space

Miki Matsumuro (muro@cog.human.nagoya-u.ac.jp)
Kazuhisa Miwa (miwa@is.nagoya-u.ac.jp)
Graduate School of Information Science, Nagoya University Furo-cho, Chikusa-ku, Nagoya, Japan

Abstract

We examined the effect of goal specificity on a search in two problem spaces: a hypothesis space and an instance space. Two hypotheses were considered: 1) a nonspecific goal facilitates a search in a hypothesis space more than a specific goal, and 2) as a hypothesis space is searched more, the performance in discovering the target rule improves. We also defined an initial hypothesis space consisting of initially considered hypotheses, and investigated the effect of this initial hypothesis space on the goal specificity effect. The results of three experiments indicated that when an initial hypothesis space was offered, the goal specificity effect was observed. A nonspecific goal actually facilitated a search in the hypothesis space. When, however, no initial hypothesis space was offered, the goal specificity effect was not confirmed. We also found that the facilitation of the hypothesis space search improved performance in discovering the target rule.

Keywords: discovery, rule induction, goal, hypothesis testing

Introduction

Dual Space Search Theory

Rule induction and scientific discovery have been studied based on the dual space search theory. Simon and Lea (1974) first suggested that a problem space consists of two spaces: a "rule space" for searching rules and an "instance space" for testing rules. Both rule and instance spaces are searched to find a correct rule.

Klahr and Dunbar (1988) extended the dual space search theory to the Scientific Discovery as Dual Search (SDDS) model for investigating scientific discovery. They considered a "hypothesis space" as a rule space and an "experiment space" as an instance space where the process of scientific discovery develops through the interaction between two types of searches in the two spaces. Reasoners state hypotheses by searching in a hypothesis space, receive feedback from an experiment space, and modify the current hypotheses or propose new hypotheses. Klahr and his colleague confirmed this model through a long series of experiments (Klahr & Dunbar, 1988; Klahr, 2000). They also identified "experimenters" who preferred to search in an experiment space and "theorists" who preferred to search in a hypothesis space. In this study we call the two spaces a "hypothesis space" and an "instance space."

The search in a hypothesis space is crucial for scientific discovery. Klahr and Dunbar (1988) demonstrated that a search in only a hypothesis space led to the discovery of a correct rule without the execution of any experiments.

Goal Specificity Effect in Dual Space Search

On the other hand, we often neglect to consider the theories or rules behind phenomena when we aim for a specific goal. We tend not to search in a hypothesis space at times like this, as we concentrate on a search in an instance space to achieve the goal.

Problem-solvers given a specific goal learn more poorly than problem-solvers given a nonspecific goal (Sweller & Levine, 1982). Burns and Vollmeyer (2002) investigated this effect of goal specificity based on the dual space search theory. Using a task in which participants were asked to learn the relations between inputs to and outputs from a system, they observed the effect of goal specificity. The NSG (nonspecific goal condition) participants, who were not informed of the target values of the outputs, learned the system structure better than the SG (specific goal condition) participants, who were informed of the target values. Burns and Vollmeyer (2002) also found that the NSG participants conducted more hypothesis testing than the SG participants. From these results, they concluded that a nonspecific goal encouraged the participants to search actively in a hypothesis space. Therefore, a nonspecific goal might lead to better learning than a specific goal.

Present Study

A hypothesis space is usually huge, hence a hypothesis space search is performed based on constraints offered by attentional perspectives (van Joolingen & de Jong, 1997). In the present study we define the hypothesis space in which the participants initially search as an "initial hypothesis space" and the space containing a target rule to be discovered as a "target space."

In the earlier studies, the initial hypothesis space was typically decided by an experimenter because the participants were informed of all the relative factors of focus. This initial hypothesis space also contained the target rule to be discovered. There was no need for the participants to find the target hypothesis space, as the initial hypothesis space and target hypothesis space were identical (Figure 1(a)). Here, in contrast, we investigate situations in which the participants must find a target hypothesis space by themselves in order to discover an appropriate rule.

(1) Initial-space situation

One situation we deal with is the "initial-space situation" (Figure 1(b)). Participants are given an initial hypothesis space by an experimenter. This initial space, however, contains no rule to be discovered. The initial hypothesis space differs from the target hypothesis space. To discover the target rule, the participants need to shift a searching space from the initial hypothesis space to the target hypothesis space. This situation typically emerges in insight problem solving (Kaplan & Simon, 1990).

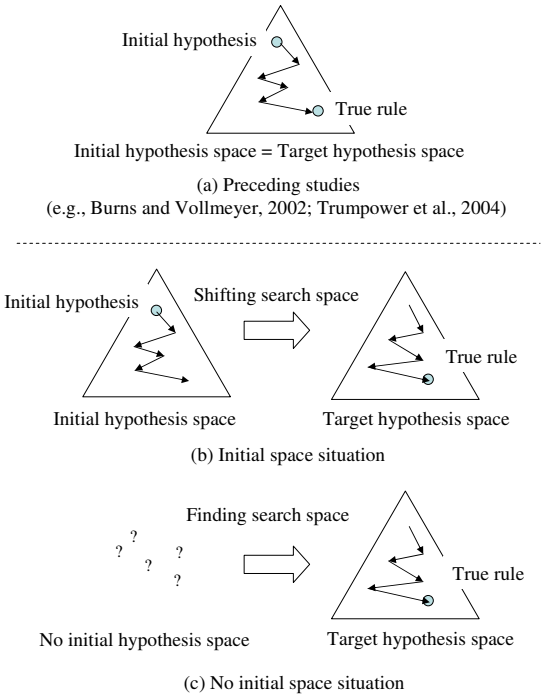

Figure 1: Conceptual diagrams of preceding and present studies

(2) No-initial-space situation

The other situation is the "no-initial-space situation" (Figure 1(c)). The participants in this case are not informed of any relative factors of focus, and thus receive no information on the initial hypothesis space to be searched. The participants have to find relative factors for the search from the initial stage. The investigation of this situation is important, as the size of a hypothesis space and the number of available hypotheses might affect the search strategies.

Aim of the Present Study

We investigate two hypotheses regarding the effect of goal specificity on a search in the dual spaces in the two situations: the initial-space situation and the no-initial-space situation.

Hypothesis 1: A nonspecific goal facilitates a search in a hypothesis space more than a specific goal. In other words, participants who are given a nonspecific goal may search more actively in a hypothesis space.

Hypothesis 2: As a hypothesis space is searched more, the performance in discovering the target rule improves.

Task

Figure 2 is a screen shot of the task for this study. The participants are asked to use the arrow buttons to pass the ball from player to player and to shoot for the basket. Two rules, one fake and one true, are valid in each game. These rules determine the relation between the arrow buttons and pass directions for the ball. In both rules, the up-arrow button corresponds to a certain direction and the other seven buttons correspond to the other seven directions relative to the up-arrow in clockwise rotation. The direction of the prior pass

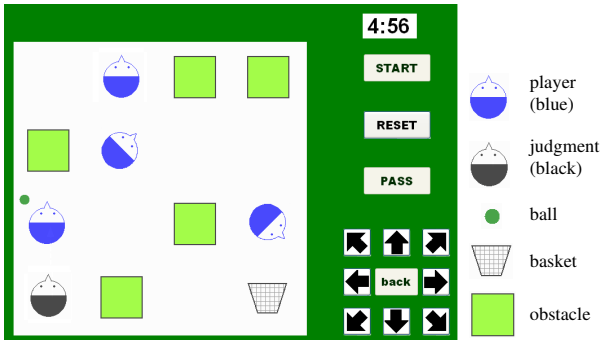

Figure 2: Screen shot of the task
A screen shot of one of the games during Phase 1. The participants pass the ball from player to player and shoot for the basket.

corresponds to the up-arrow button in the true rule, whereas the direction of the current player's nose corresponds to the up-arrow button in the fake rule. Although the fake rule is expected to be discovered more easily than the true rule, it can be valid in the initial games (not in all games). The true rule, in contrast, is valid in all games.

The participants, having discovered the fake rule, initially search in the hypothesis space consisting of hypotheses characterized by a face direction (the "face hypothesis space"). Even if the fake rule no longer validly works in the games, the participants continue to search in the face hypothesis space. To discover the true rule, the participants must shift a searching hypothesis space from the face hypothesis space to the hypothesis space made up of hypotheses characterized by the orbit of the prior pass (the "orbit hypothesis space"). The orbit hypothesis space contains the true rule. Thus, in this task, the initial hypothesis space means the face hypothesis space and the target hypothesis space means the orbit hypothesis space.

The experiment basically consists of three phases.

Phase 1: The participants engage in games in which both the fake and true rules are valid. The participants are asked to shoot the ball into the basket as many times as possible. We expect the participants to discover the fake rule and use it in Phase 1.

Phase 2: Only the true rule can be applied in the games in Phase 2. At the beginning of Phase 2, the participants are expected to search in the face hypothesis space, based on their successes with the fake rule accumulated in Phase 1. To discover the true rule, the participants must shift a searching hypothesis space from the face hypothesis space to the orbit hypothesis space. We manipulate experimental factors and observe how these factors affect the searches in the hypothesis space. The playing time is limited in Phases 1 and 2. The participants are told that the games in these phases are for training, and that the real games, or the actual takes, will come in Phase 3.

Phase 3: The participants are informed the real games, or actual takes, come in Phase 3. They are asked to shoot the ball into the basket. Phase 3 consists of two games, each of which

is played to completion without a time limit imposed. The true rule is valid in both the first and second games, but the fake rule can be applied to only the first game. We, the researchers, judge whether or not the participants discover the true rule based on their performances in each game.

Manipulation of Goal Specificity

Goal specificity is manipulated mainly through the following three experiments. The participants are given a specific goal (the SG condition) and asked to play the games in Phase 2 (the screen shot in Figure 2 shows a Phase 2 game). The basket and obstacles determine only one pass route. Thus, the participants' next moves are specified. Meanwhile, the participants in the other group are given a nonspecific goal (the NSG condition) and asked to play games in which there are no obstacles and in which the basket is replaced by a player. In this situation, the next move is unspecified: a participant can intentionally select one of several valid passes without a specific final goal (basket). With these manipulations in the games come differences between the instructions under the SG and NSG conditions in Phase 2. The participants in the SG condition are asked to shoot the ball into the basket, whereas the participants in the NSG condition are asked to pass the ball from player to player. In both conditions, the participants are asked to perform as many games as possible within the time limit. The time point of every button selection by a participant is recorded.

Measurement The hypothesis space in this task consists of hypothesized rules on the relations between the arrow buttons and the pass directions. An instance space consists of all instances; each instance is described as "when a certain arrow button is selected, a ball is thrown to a certain direction under a certain situation." Assuming that a search in one space is performed after a search in the other space, in turn, a hypothesis space search is performed during the period elapsed between the receipt of one pass result (the result of one pass thrown) to the receipt of the next pass result. Therefore, in this study, we use a time interval of two successive passes as a measurement for the amount of searches in a hypothesis space. Henceforth we refer to this time interval as the "pass interval time."

We judged whether each participant discovered the true rule from his or her performance in Phase 3. If the participants could not discover the true rule, the adjustment strategy minimized errors. The participants who use the adjustment strategy make a pass at first based on some criterion or randomly, and then adjust the direction of arrow buttons in order to minimize the difference between the expected and actual pass directions. We defined the successful participants as the participants whose error rate was lower than the expected error rate when they use the adjustment strategy.

Experiment 1

We conducted Experiment 1 to investigate the effect of goal specificity on a search in a hypothesis space in the initial-space situation (see Figure 1(b)). In addition, we manipulated another factor, the instruction factor, to test whether the pass interval time is valid as a measurement of the amount of searches in a hypothesis space. In the search-oriented condition (the SO condition), the participants were asked to find a rule that determines the relation between the arrow buttons and pass directions. By contrast, in the non-search-oriented condition (the NSO condition), the participants were told nothing about the rule. This manipulation may lead the participants in the SO condition to search more in the hypothesis space, compared to the participants in the NSO condition. If the pass interval time correlates with the amount of searches in the hypothesis space, the pass interval time of the participants in the SO condition will exceed that of the participants in the NSO condition.

Method

Participants Sixty-four undergraduates participated in Experiment 1. Each was assigned to one of four conditions: goal specificity (SG and NSG) $\times$ instruction (SO and NSO).

Task and Procedure Experiment 1 was conducted in small groups of three or fewer participants. After the participants received a basic explanation of the procedures, the participants briefly rehearsed the task. Next, they carried out the task in the three phases. Phase 1 and Phase 2 each lasted for five minutes. In Phase 2, two factors: the participants' search preferences in the hypothesis space by the instruction and goal specificity, were manipulated. Finally, in Phase 3, all participants played two games in an identical situation without a time limit imposed.

Results and Discussion

Pass Interval Time Figure 3 presents the average pass interval time in each condition in Phase 2. A two-way ANOVA ((goal specificity: SG and NSG) $\times$ (instruction: SO and NSO)) was performed on the pass interval times in Phase 2. The interaction between the two factors was not significant ($F(1,60) = 0.673, n.s.$). The main effects of both the goal specificity factor ($F(1,60) = 38.454, p < .001$) and the instruction factor ($F(1,60) = 5.030, p < .05$) reached significance.

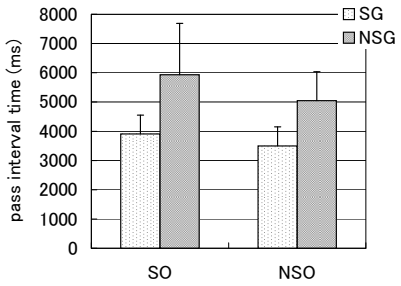

Figure 3: Average pass interval time in each condition in Phase 2 in Experiment 1 (bars show standard errors)

For the instruction factor, the pass interval time of the participants in the SO condition was longer than that of the participants in the NSO condition. The instruction given in the SO condition, the instruction which encouraged the participants to find a rule, increased the pass interval time. Noting

that this implied a correlation between the pass interval time and the amount of searches in a hypothesis space, we decided to use the pass interval time as a measurement of the amount of searches in a hypothesis space.

For the goal specificity factor, the pass interval time of the participants in the NSG condition was longer than that of the participants in the SG condition. This corroborated the first hypothesis: a nonspecific goal facilitates a search in a hypothesis space.

Proportion of Successful Participants Next, we analyzed the proportion of participants who discovered the true rule in each condition. For the instruction factor, 6 of 33 participants in the SO condition and 2 of 31 participants in the NSO condition were successful. There was no significant difference between the two conditions ($p > .10$). Similarly, for the goal specificity factor, 4 of 30 participants in the SG condition and 4 of 34 participants in the NSG condition were successful. Again, there was no significant difference between the two conditions ($p > .10$). Hence, these results did not confirm the second hypothesis: more searches in a hypothesis space improve performance in discovering the target rule.

Experiment 2

In Experiment 2 we investigated the effect of goal specificity on a search in a hypothesis space in the no-initial-space situation (see Figure 1(c)). We also manipulated the instruction factor to test the validity of the pass interval time, as was done Experiment 1.

Method

Participants Sixty-four undergraduates participated in Experiment 2. Each was assigned to one of four conditions: goal specificity (SG and NSG) $\times$ instruction (SO and NSO).

Task and Procedure The task in Experiment 2 was almost the same as that in Experiment 1, with the following adjustments. No Phase 1 was conducted in Experiment 2. The faces were removed from the players and the referee, and replaced with blue- and black-filled circles. The participants did not acquire the initial hypothesis space, as they were given no perspectives on which to focus for forming hypotheses at the beginning of the task.

Results and Discussion

Pass Interval Time Figure 4 presents the average pass interval time in each condition in Phase 2. A two-way ANOVA ((goal specificity: SG and NSG) $\times$ (instruction: SO and NSO)) was performed on the pass interval times in Phase 2. The interaction between the two factors was not significant ($F(1,60) = 0.022, n.s.$). The main effects of both the goal specificity factor ($F(1,60) = 6.708, p < .05$) and the instruction factor ($F(1,60) = 4.056, p < .05$) reached significance.

In the analysis for the instruction factor, this result was consistent with the result in Experiment 1. The pass interval time of the participants in the SO condition was significantly longer. The correlation between the pass interval time and the searches in a hypothesis space was again supported.

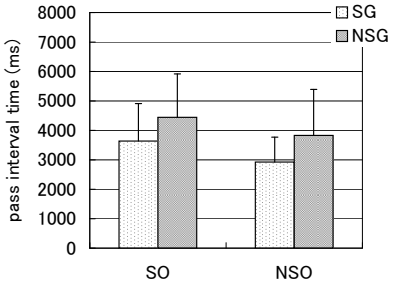

Figure 4: Average pass interval time in each condition in Phase 2 in Experiment 2 (bars show standard errors)

In the analysis for the goal specificity factor, the pass interval time was longer in the NSG condition than in the SG condition. This result also confirms the first hypothesis, corroborating the finding of Experiment 1. Note, however, that the difference between the SG and NSG conditions was much smaller in Experiment 2 than in Experiment 1. We will discuss this difference further in a later section.

Proportion of Successful Participants Next, we analyzed the proportion of participants who discovered the true rule in each condition. In the analysis for the instruction factor, 4 of 32 participants in the SO condition and 5 of 32 participants in the NSO condition were successful. There was no significant difference between the two conditions ($p > .10$). The result was similar in the analysis for the goal specificity: 3 of 32 participants in the SG condition and 6 of 32 participants in the NSG condition were successful. Again, there was no significant difference between the two conditions ($p > .10$). Hence, these results did not confirm the second hypothesis.

Comparison between Experiments 1 and 2

By comparing the results of Experiments 1 and 2, we could explore how the existence of the initial hypothesis space affected the effect of goal specificity. A premise for this study was dual spaces for search: the hypothesis space and instance space. Yet the participants in the NSO condition may not have assumed any hypothesis space, as they may not have noticed the rule determining the pass directions. For this reason, the following analysis focused on the participants in the SO condition.

In these experiments we introduced what we called the "situation factor," manipulating whether or not the participants had the initial hypothesis space across Experiments 1 and 2. At the beginning of Phase 2, the participants in Experiment 1 had the initial hypothesis space. Recollecting their accumulated successful experiences with the fake rule in Phase 1, they directed their attention to the face hypothesis space. This situation was called the "initial-space condition" (the IS condition). In contrast, the participants in Experiment 2 did not acquire an initial hypothesis space or experience any game play in Phase 1. And by removing the faces as cues from the players of the games in Phase 2, we deprived the participants of perspectives for forming hypotheses. This situation was called the "no-initial-space condition" (the NIS condition).

A two-way ANOVA ((situation: IS and NIS) $\times$ (goal speci-

ficity: SG and NSG)) was performed on the pass interval times in Phase 2. As a result, a marginally significant interaction between the situation and goal specificity factors was revealed ($F(1,61) = 3.158, p = .081$). In the IS condition, the pass interval time was longer in the NSG condition than in the SG condition ($F(1,61) = 17.449, p < .001$). This effect, however, disappeared in the NIS condition ($F(1,61) = 2.769, n.s.$). Both the goal specificity and situation factors had significant effects ($ps < .05$).

In this comparison, the participants with a nonspecific goal had a longer pass interval time than the participants with a specific goal in the IS condition. Meanwhile, the goal specificity factor had no effect on the pass interval time in the NIS condition. Thus, the first hypothesis was confirmed only in the IS condition, and not in the NIS condition.

Experiment 3

In Experiments 1 and 2, we found that goal specificity had no effect on a search in a hypothesis space when the participants lacked an initial hypothesis space. In Experiment 3, we manipulated both the goal specificity and situation factors to confirm the effect of these factors directly. Several of the experimental procedures were improved for this experiment. First, Phase 1 was performed in both the IS and NIS conditions, so that the participants would begin Phase 2 with identical prior experiences. In the NIS condition, the faces of the players and referee were removed in Phase 2 to eliminate the initial hypothesis space. Second, only a few participants successfully discovered the target rule in Experiments 1 and 2. In Experiment 3, the players who threw a successful pass and the receiver from the previous trial were marked visually on the game display. This cue lowered the memory loads of the participants, thus helping the participants discover the true rule in the orbit hypothesis space more easily.

Method

Participants Seventy-four undergraduates participated in Experiment 3. Each was assigned to one of four conditions: situation (IS and NIS) $\times$ goal specificity (SG and NSG).

Task and Procedure Experiment 3 was conducted in small groups of three or fewer participants. To control prior experiences, the participants in all conditions played games in all three phases. In Phase 1, the participants played games in which both the fake and true rules were valid, over a total play time of five minutes. In Phase 2, the participants played games in which only the true rule was valid. The goal specificity factor was manipulated by the same method used in the prior two experiments. Additionally, the situation factor was manipulated by adjusting the players' faces. The participants in the IS condition played the games with normal face players, as they had in Experiment 1. Meanwhile, the participants in the NIS condition played the games with faceless players, as they had in Experiment 2. Unlike Experiments 1 and 2, the game time in Phase 2 was increased to seven minutes in order to increase the number of successful participants. All participants were instructed that there was a rule valid through all of the games. Finally, Phase 3 was conducted using the same player faces used in Phase 2, but without a time limit.

Results and Discussion

Pass Interval Time Figure 5 presents the average pass interval time in each condition in Phase 2. A two-way ANOVA ((situation: IS and NIS) $\times$ (goal specificity: SG and NSG)) was performed on the pass interval times in Phase 2. The interaction between the situation and goal specificity factors reached significance ($F(1,70) = 4.989, p < .05$). In the IS condition, the pass interval time was longer in the NSG condition than in the SG condition ($F(1,70) = 9.078, p < .005$). This effect disappeared, however, in the NIS condition ($F(1,70) = 0.021, n.s.$). The goal specificity and situation factors both had significant effects ($ps < .05$).

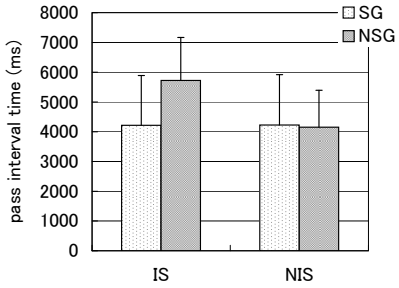

Figure 5: Average pass interval time in each condition in Phase 2 in Experiment 3 (bars are standard errors)

The result of Experiment 3 was consistent with the results of Experiments 1 and 2. In the IS condition, where the participants acquired the initial hypothesis space, goal specificity had an observable effect on a search in a hypothesis space. The participants with a nonspecific goal searched in a hypothesis space more actively than the participants with a specific goal. This effect was not observed, however, in the NIS condition, where the initial hypothesis space was eliminated by the change of the game display. Therefore, the presence or absence of an initial hypothesis space affected the goal specificity effect in a search in a hypothesis space. The first hypothesis is confirmed only in the IS condition.

Proportion of Successful Participants Next, we analyzed the proportion of participants who discovered the true rule in each condition (Figure 6). In the IS condition, 2 participants discovered the true rule in the SG condition and 8 participants discovered the true rule in the NSG condition.

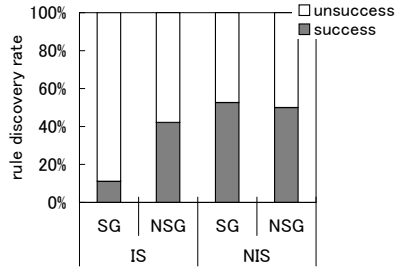

Figure 6: Proportion of successful participants in each condition in Experiment 3

The participants in the NSG condition discovered the true rule more frequently than the participants in the SG condition ($\chi^2(1) = 4.502, p < .05$). In the NIS condition, 10 participants discovered the true rule in the SG condition and 9 participants discovered the true rule in the NSG condition. There was no significant difference between the SG and NSG conditions in the NIS condition ($p > .10$).

In the IS condition, the pass interval time of the participants in the NSG condition was longer than that in the SG condition. Additionally, the proportion of successful participants in the NSG condition was also greater than that in the SG condition. Thus, the second hypothesis is confirmed.

Discussion and Conclusions

In this study we investigated the following two hypotheses in two situations, an initial-space situation and a no-initial-space situation: (1) A nonspecific goal facilitates a search in a hypothesis space rather than a specific goal. (2) As a search in a hypothesis space is more actively performed, the performance in discovering the target rule improves.

From the results of three experiments, the first hypothesis was partially confirmed. The effect of goal specificity on a search in a hypothesis space depended on whether or not the participants noticed an initial hypothesis space. When the participants noticed an initial hypothesis space, goal specificity had an observable effect on a search in a hypothesis space. The participants with a nonspecific goal searched in a hypothesis space more actively than the participants with a specific goal. On the other hand, this effect of goal specificity was not observed when the participants were not given any initial hypothesis space.

The second hypothesis was confirmed in the results of Experiment 3. The pass interval time of the participants with a nonspecific goal was longer than that of the participants with a specific goal in the IS condition. Additionally, the proportion of successful participants with a nonspecific goal was larger than that of successful participants with a specific goal in the IS condition. However, there was no significant difference between the specific goal and nonspecific goal conditions in Experiments 1 and 2. This may have been due to a floor effect, as only a few participants discovered the true rule in these experiments.

The results of the present study are consistent with the finding of Klahr and Dunbar (1988). They defined hypotheses as the forms of a "frame." In their study, they classified the frames (hypotheses) into several types, according to their features. The hypothesis spaces in our study, i.e., sets of hypotheses sharing a common feature, could be explained by the types of frames defined by Klahr and Dunbar (1988). The theorists in Klahr's experiments preferred to do their searches in hypothesis spaces. They were able to switch the hypotheses types correctly, within short periods of time and over the course of only a few experiments, and discovered the rule rapidly. In our study, the participants with nonspecific goals behaved like the theorists in the situation where the initial hypothesis space was given. They preferred to search in a hypothesis space, repeating the behavior of the theorists in the earlier studies. They were able to switch the searching hypothesis space from a given initial hypothesis space to a target hypothesis space with fewer instances, and discovered the true rule. In contrast, the participants with a specific goal preferred to search in an instance space, repeating the behavior of the experimenters defined in Klahr's study.

In the initial-space situation we observed the effect of goal specificity on a search in a hypothesis space, duplicating the results from earlier studies. This situation is identical to situations covered in the preceding studies, where participants were given an initial hypothesis space. Unlike the preceding studies, we used a task in which the true rule was not included in the initial hypothesis space. To discover the true rule, the participants needed to shift their attention to the target hypothesis space. Here, the effect of goal specificity on a search in a hypothesis space was still confirmed.

Yet when the participants were given no initial hypothesis space, goal specificity had no observable effect on a search in a hypothesis space. To state hypotheses, the participants initially needed to find a focused hypothesis space by themselves in this situation. We assume that they searched in an instance space to collect data as cues for determining a hypothesis space to search. This may explain why the hypothesis space search was not activated for the participants with the nonspecific goal. The SDDS model proposed that a discovery process is controlled with three main components: "search hypothesis space," "test hypothesis," and "evaluate evidence." The search hypothesis space component corresponds to a search in a hypothesis space in our study. This component contains a search in an experiment (instance) space as one of the sub lower components. Participants could collect data and find a pattern of these data gathered through the experiment space search, and state hypotheses. Similarly, the participants in our study who were given no initial hypothesis space needed cues to find a focused hypothesis space and state hypotheses. Therefore, we conclude that they searched in an instance space regardless of goal specificity.

References

Burns, B. D., & Vollmeyer, R. (2002). Goal specificity effects on hypothesis testing in problem solving. *The Quarterly Journal of Experimental Psychology, 55A*, 241–261.

Kaplan, C. A., & Simon, H. A. (1990). In search of insight. *Cognitive Psychology, 22*, 374–419.

Klahr, D. (2000). *Exploring science*. The MIT Press.

Klahr, D., & Dunbar, K. (1988). Dual space search during scientific discovery. *Cognitive Science, 12*, 1–48.

Simon, H. A., & Lea, G. (1974). Problem solving and rule induction: A unified view. In I. L. W. Gregg (Ed.), *Knowledge and cognition.* Potomac, Md: L. Erlbaum Associates.

Sweller, J., & Levine, M. (1982). Effects of goal specificity on means-ends analysis and learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8*, 463–474.

van Joolingen, W. R., & de Jong, T. (1997). An extended dual search space model of scientific discovery learning. *Instructional Science, 25*, 307–346.