

How Expert Tutors Revise Tutoring Policies and Strategies When Students Make Mistakes

Evelyn Lulis (elulis@cdm.depaul.edu)

CDM, DePaul University
243 S. Wabash Avenue, Chicago, IL 60604 USA

Shlomo Argamon (argamon@iit.edu)

Martha Evens (evens@iit.edu)

Department of Computer Science, Illinois Institute of Technology
10 West 31st Street, Chicago, IL 60616 USA

Abstract

This paper shows that expert tutors adjust both their high-level policies and their lower-level strategies, based on assessments of student performance. This contrasts with previous reports that tutors do not use such sophisticated methods in their tutoring. The specific strategy adjustments that we find are consistent with the interpretation that expert tutors seek to aid weaker students in practicing problem solving skills, perhaps to build a new mental model of the domain, and seek to aid stronger students in diagnosing and correcting errors in their current mental model. These results have implications for better understanding of human tutoring, as well as for the design of effective intelligent tutoring systems. We believe that computer tutors should be designed to use student assessments to determine **how** they tutor as well as what they tutor.

Introduction

Previous studies on tutoring have claimed that application of different tutoring strategies is rarely directed towards diagnosing student modeling errors (Graesser, Person, & Magliano, 1995), and that a tutor's accurate assessment of the student's capabilities has been deemed largely irrelevant to tutoring effectiveness (Chi, Siler, & Jeong, 2004). Intelligent tutoring systems that have built cognitive models of the student have utilized their models to determine **what** subject areas to focus on in a tutorial session, but have not adjusted **how** they tutor. However, recent results have shown that expert tutors dynamically adjust their tutoring policies in response to changes in their assessments of student abilities (Cho, 2000; Cho, Michael, Rovick, & Evens, 2000); previous contradictory results may be explained by the fact that subjects in those cases were novice tutors who were experts in their subject, but inexperienced in tutoring.

Expert tutors plan constantly before and during tutoring sessions based on their previous experience with this and other students (Evens & Michael, 2006; Leinhardt & Greeno, 1986). In this paper, we distinguish between three levels of plans: high-level *policies*, which define the overall structure of the tutoring session, lower-level *strategies*, which are chosen dynamically during a session, to accomplish specific educational subgoals, and *tactics*, such as probing or hinting, which are used in many different

strategies (discussed in Kim, Freedman, Glass, & Evens, 2006). The strategies used in our analysis include summarizing, explaining via analogy, moving forward in a chain of causal reasoning, and so forth. Strategies typically involve several turns in the dialogue, while tactics usually involve just one. Policies are usually set for a whole stage in a problem, or even a whole problem, or a whole session. We set out to analyze how tutors select different strategies to implement different policies for different students with the goal of increasing our understanding of how tutors adapt their tutoring strategies as well as their choice of material based on student assessment.

The research described in this paper originated in an accidental discovery by Byung-In Cho, when he was attempting to model tutor assessments of student performance (Cho, 2000; Cho, Michael, Rovick, & Evens, 2000). Cho set up several modeling criteria including success in the current discussion (more than half of the student's responses to tutor questions are correct) and the percentage of correct predictions made by the student in the current stage. He decided to test these criteria using a set of nine two-hour long tutoring sessions carried out in 1992, just at the point where the expert tutors had defined a new tutoring policy for Version 2 of CIRCSIM-Tutor. In the original *Immediate Feedback (IMF) policy*, the tutor helps the student solve the problem step by step, commenting immediately on every student input, good or bad. In the new *V2 policy*, the tutor attempts first to build up a model of the student's understanding, by asking the student to predict the qualitative changes in all the important variables in the model of the baroreceptor reflex, and then to gear tutoring to correcting any misconceptions that the tutor diagnoses. Michael and Rovick planned the series of human tutoring sessions studied here to provide us with samples of tutoring language using the V2 policy for use in building the system.

Cho observed that the tutors, in spite of their commitment to provide examples of the V2 policy, sometimes abandoned the V2 policy in favor of the IMF policy. We had never noticed these policy switches and were startled to see them. With the help of the machine learning program, C5.0, Cho was able to establish that switches from V2 to IMF occurred when and only when the student failed to achieve success in discussion, made repeated errors in determining the primary variable, or made more than two prediction errors in the

current phase. Switches from IMF back to V2 occurred at the beginning of a new phase after the student started to perform better.

It was decided that if we were going to implement this kind of adaptation we needed to study the strategies used by our human tutors and discover how the choice of a tutoring policy determined the choice of strategies used.

The expert tutors were Joel Michael and Allen Rovick, Professors of Physiology at Rush Medical College; the students were first-year medical students who were typically highly competent at memorizing new material, but often experienced difficulty when they had to apply this information to new situations; they therefore needed to practice using qualitative reasoning to solve problems. The tutors presented the student with a problem, a perturbation in the blood pressure, and instructed them to enter predictions regarding qualitative changes in seven important cardiovascular variables in a prediction table (Rovick & Michael, 1992). After the predictions were entered, the tutor began a Socratic tutoring dialogue with the student in natural language, to help remedy prediction errors.

The overall tutoring process was divided into three stages corresponding to the three stages in the blood pressure stabilization process: the Direct Response (DR) period before the reflex kicks in, the Reflex Response (RR), and the new Steady State (SS) attained after a few minutes. In addition, the tutor made certain that the student began by correctly identifying and predicting the behavior of the primary variable, the first variable to change in the DR phase. A mistake at this point could result in the propagation of errors throughout the Prediction Table. A series of experiments with CIRCSIM-Tutor (Evens & Michael, 2006; Michael, Rovick, Glass, Zhou, & Evens, 2003) demonstrated that using the system for an hour produces learning gains significantly greater than those produced by reading targeted text for an hour (cf. VanLehn et al., 2007). This system is now in routine use at Rush Medical College.

The strategies identified in the sessions studied are similar to those identified by other researchers (Fox, 1993; Chi, Siler, Jeong, Yamauchi, & Hausmann, 2001; Graesser, Person, & Magliano, 1995). But like Fox we have found some relatively sophisticated strategies.

- **T-does-neural-DLR (DLR)** A directed line of reasoning explaining why neural variables do not change in DR.
- **T-tutors-via-determinants (DET)** An interactive exposition of a fundamental causal reasoning process based on analysis of determinants of the current variable.
- **T-moves-forward (MVF)** The tutor asks for the next step in the reasoning process.
- **T-prompts-start (STA)** Gives the student a boost at figuring out what to predict first.

- **T-explores-anomaly (EXA)** Helps the student understand what happens when two determinants move in opposite directions.
- **T-shows-contradiction (CON)** Asks the student a series of questions that demonstrate that his or her prediction is impossible.
- **T-diagnoses-error (DIA)** Asks the student to explain his or her reasoning and then debugs the response.
- **T-tutors-via-analogy (ANA)** Proposes a helpful analogy and asks the student to make an inference based on it.
- **T-tutors-algebraic approach (ALG)** Explains how to make SS predictions by computing a logical sum of DR and RR predictions.
- **T-tutors-via-deeper-concepts (DPR)** This strategy discusses the biochemistry or anatomy underlying a change.
- **T-moves-toward-PT (PRE)** Forces the student to think about the implications of the current change for variables in the Prediction Table.
- **T-tutors-via-negative-reflex (NGR)** Discusses the cause and effect relationships involved in the baroreceptor reflex.
- **T-summarizes (SUM)** Summarizes the current state of the variables under discussion or gets the student to do so.
- **T-tutors-logical-order (LOG)** Helps the student figure out what happens to each determinant before predicting a new result.
- **T-probes-correct-answer (PRO)** Begins with a probe just like T-diagnoses-error, but ends with an assurance to the students that inputs will be handled carefully.
- **T-tutors-the-rules-of-the-game (ROG)** Explains assumptions about the domain model and about the tutoring model.
- **T-tutors-misconception (MIS)** Tries to provide an alternative interactive explanation.

More detailed descriptions of all but the last three of these strategies can be found in (Kim, Freedman, Glass, & Evens, in press).

Hypotheses

As we started to examine the distribution of strategies in different policies we began to develop hypotheses about the underlying tutor choices.

It was hypothesized that, for weaker students, tutors would tend to choose to assist them in creating a new domain model from scratch, either because the student lacks one, or because the student models are too far off the mark for reasonable diagnosis or correction to be possible. For stronger students, tutors seek to diagnose more localized bugs in student domain models, and help students correct such bugs in a focused manner. We therefore anticipated that:

1. The IMF policy, preferred by expert tutors for weaker students, would focus on helping the

student practice solving basic problems; hence, strategies used in IMF would be those that quickly moved the student through problem solving: Prompts-Start, Moves-Forward, Logical-Order, and Tutors-via-Determinants. As well, we expected to see more strategies overall in each tutoring phase.

2. The V2 policy, preferred for stronger students, would employ fewer problem solving strategies, and instead would use more diagnostic and model-corrective strategies, such as Tutors-via-Analogy, Shows-Contradiction, or Probes-Correct-Answer.

Data and Methodology

The data examined in this study came from nine tutoring sessions carried out in November 1992. The tutors were Allen Rovick and Joel Michael; the students were volunteers from their first-year class in Human Physiology. Tutor and student were typing on computers in two separate rooms communicating through telephone modems. The interaction was recorded on the hard drive by CDS, the Computer Dialogue system, written by Jun Li (Li, Seu, Evens, Michael, & Rovick, 1992). CDS enforced turn taking, but allowed one party to interrupt the other. The nine transcripts are numbered K30-K38.

Rovick served as the tutor in five sessions, K30-K34, and Michael in the other four, K35-K38. The tutors taught the same two problems in each session, but alternated them so that the problem that was second in one session was first in the next. Typically, each problem was taught in three phases, DR, RR, and SS, but sometimes the tutor skipped phases for lack of time. Although the tutors produced these tutoring sessions largely so that we could analyze them and use the results in our research, they were not satisfied themselves with their own primary variable tutoring and suggested that we not use that portion of the data. As a result, portions of the sessions that involve primary variable tutoring were not used in this study. We also excluded student initiatives and tutor response strategies, which were studied in (Shah, Evens, Michael, & Rovick, 2002). Thus the dialogue studied is entirely tutor directed.

Our methodology was simply to mark up transcripts from human tutoring sessions in SGML and then count the strategies used in each phase. This technique was used by Pilkington (1999) and by the Dialogue Research Group under the leadership of Johanna Moore and James Allen (Allen & Core, 1997). A previous study of tutorial dialogues involving the same two tutors was based primarily upon sessions that follow the V2 policy (Kim et al., 2006).

Two of the authors classified and marked up the policies and strategies in nine keyboard tutoring sessions, K30-K38. The first eight keyboard sessions, K1-K8, using the IMF policy, were used as training material. K1 was jointly marked up and a list of the strategies found compiled. K2-K4 were marked up separately, with only a marginal level of agreement. Sessions K5-K8 were searched for further examples of these strategies and the borderline cases discussed extensively. All but three of the strategies

discovered appear in Kim et al. (2006). A manual was written, adding three strategies to the ones described by Kim, with definitions and examples to assist in the markup process.

Most of the strategies on our list are described in detail by Kim et al. (2006), but three of those found in the sessions studied do not appear in that paper: T-probes-correct-answer, T-tutors-logical-order, and T-tutors-ROG. The T-probes-correct-answer strategy is very much like the T-diagnoses-error strategy. In both cases, the tutor starts by asking the student how a particular prediction was made, or why the student made this prediction. The difference appears at the end; in the T-diagnoses-error strategy the tutors asks for a new prediction. It was expected that T-tutors-logical order would be much more common in the IMF policy than in the V2 policy, because there is much more emphasis on discussing the complete problem with the student. Finally, it was thought that the T-tutors-rules-of-the-game strategy might be more likely to occur with the IMF policy, because the weaker students might experience more difficulty understanding the rules of the game. Also included in the Rules-of-the-Game category were requests that the student reread the problem description, or other reminders to consider the original problem.

Work on K30-K38 began by dividing the sessions into phases and the phases into separate tutoring episodes. Without mentioning the strategy names, the episodes were divided into separate territory for each strategy before starting the classification process. The strategies in all nine sessions were marked up separately, with agreement on 235 and disagreement on 10. Computation of Cohen's kappa (Carletta 1996; Di Eugenio 2000), showed significant inter-rater reliability with kappa=0.95.

Results

Full results are given in Table 3; statistical significance was tested, where appropriate, using a one-sided t-test assuming unequal variances. Examining differences in individual strategies, our hypothesis that the Move-Forward strategy would be more frequent under the Immediate Feedback (IMF) policy was supported. There were 75 in the 15 IMF phases (a mean of 5 per phase), but only 15 in the 31 V2 phases (giving a mean of 0.48); this difference is highly significant with $p < 0.0001$. There were five examples of T-prompts-start under IMF, but none under V2; this difference is significant with $p < 0.05$. Finally, there were 6 examples of T-tutors-logical-order under IMF and only 3 under V2. The t-test gives significance with $p < 0.05$, but with this small N, this conclusion is merely suggestive. On the other hand, the hypothesis that tutoring via determinants would be more common under IMF policy was incorrect, however—there are 7 examples under V2 and 5 under IMF. We believe that this is because the tutors generally favor this strategy because it emphasizes the style of causal reasoning that they want students to internalize, and so it appears in many sessions.

Many of the other strategies seem to be evenly distributed between the two policies. Diagnosing an error turned out to

be the most frequent strategy after moving forward with 26 examples split 15 to 11. Of course, this is something the tutor needs to do under any policy. The third most common strategy was summarizing, with 18 examples split 10 to 8. The fourth most frequent strategy was tutoring via the negative reflex, a core concept here, with 12 examples, split 7 to 5. Exploring anomalies was split 3 to 3. So was returning to the prediction table. Various approaches to tutoring the problem-solving algorithm seem to be divided fairly evenly, such as tutoring in a logical order and tutoring the algebraic approach to solving SS.

As expected, though, several of the strategies that require more sophisticated understanding from the student are more frequent under V2. There were 11 examples of tutoring via analogy under V2 and only 2 under IMF. Applying Fisher's Exact Test to the matrix in Table 1 demonstrates a significant difference with $p < 0.05$.

Table 1: Occurrence of the Tutors-via Analogy Strategy

Strategies	V2	IMF	TOTALS
ANALOGY	11	2	13
OTHER STRATEGIES	88	134	222
TOTALS	99	136	235

There were six examples of probing a correct answer under V2 and three under IMF. Fisher's Exact Test on the matrix in Table 2 gives a possibly suggestive p of 0.089.

Table 2: Occurrence of the Probes-Correct-Answer Strategy

Strategies	V2	IMF	TOTALS
PROBE CORRECT ANS.	6	3	9
OTHER STRATEGIES	93	133	226
TOTALS	99	136	235

We eliminated the three strategies with the lowest N from Table 3 to save space (thus the total shown there is 225). There were four examples of tutoring by contradiction under the V2 policy and none under the IMF policy. The binomial test and Fisher's Exact Test both give a probability of 1/16 or 0.0625 for this outcome. The strategies tutoring via deeper concepts and tutoring a misconception each occurred only three times, but those three times were all in V2 sessions (one each in three different sessions for both of them). This pattern of occurrence has a probability of 1/8 or 0.125, so these results are only suggestive, but taken in combination they suggest that the tutors deliberately choose to use these strategies with the higher-performing students.

Turning to an overall look at the immediate feedback policy, it was observed that five of the sessions involved policy breaks from the V2 policy to the IMF policy and four did not. As its definition implies, tutoring with IMF takes longer. It is therefore unsurprising that all of the occasions when material was skipped are associated with IMF. Furthermore, many more tutoring strategies were deployed when this policy was in use. The average number was 9.06

per tutoring phase vs. 3.19 per phase for the V2 policy and a t-test showed that this difference was significant at the 0.0001 level ($p = 3.57 \times 10^{-8}$).

Discussion

As noted above, some previous studies (Graesser, Person, & Magliano, 1995; Chi, Siler, & Jeong, 2004) did not find student assessment or sophisticated tutoring strategies useful in tutoring. In contrast, during the 18 hours of tutoring examined here, many sophisticated strategies were deployed by expert tutors, who tuned both policy and strategy on the basis of changing student assessments. We believe that previous studies did not see such phenomena because they employed novice tutors, rather than experts. Glass, Kim, Evens, Michael, & Rovick (1999) have found a number of striking differences between expert and novice tutors. Novice tutors produce much longer turns; experts make their students talk instead. Novice tutors give their students the answer; expert tutors make the students do the work. Novice tutors introduce extraneous concepts; experts focus on the problem.

The results we present here show how expert tutors dynamically switch their tutoring policies and strategies to match their assessments of student abilities. One of our next steps is to translate these results into plans in Freedman's (2000a, b; 2001) APE planner and add them to the repertoire of plans already in CIRCSIM-Tutor. After implementing both policies together with policy-switching methods, we can examine which policy, or combinations of policies, is most effective for which sorts of students. As well, we can investigate the development of more sophisticated tutoring policies and switching mechanisms.

Further data analysis is also crucial; a couple of the results in this paper are borderline significant, and analysis of more tutoring sessions will enable us to further elucidate any effects. Such analysis will also enable us to address another research goal—examining whether tutoring has an effect on student language as well as on student problem-solving. Do tutoring dialogues succeed in teaching the sublanguage of physiology, or the sublanguage of causal reasoning? We have never carried out the necessary study, but examining these sessions has made us realize that they contain much evidence of language teaching, as the tutors rephrase student answers in a more domain-appropriate style and call for conversational repair. Such a study would have implications both for the design of more effective ITS as well as possibly for training more effective human tutors.

Conclusions

These results clearly show how expert tutors change both their tutoring policy as well as their lower-level tutoring strategies, in response to changing assessments of student competence. For students assessed to be weaker, tutors switched to a policy and strategies designed to drill students in problem solving, ignoring any mental domain models that students may have developed. For students assessed to be stronger, tutors used a policy and strategies designed to

probe and correct students' mental models. Such strategies generally require more sophistication from the student, and so are more appropriate for higher performing students. We believe that these results have implications for the use of student modeling in ITS, in that such models should be used to modulate **how** tutoring is done, not just **what** material is to be focused on.

Acknowledgments

This work was partially supported by the Cognitive Science Program, Office of Naval Research under Grants No. N00014-94-1-0338 and N00014-02-1-0442 to Illinois Institute of Technology. The content does not reflect the position or policy of the government and no official endorsement should be inferred.

References

- Allen, J., & Core, M. (1997). Draft of DAMSL: Dialog act markup in several layers. Available from the University of Rochester, Department of Computer Science at <ftp://ftp.cs.rochester.edu/pub/packages/dialog-annotation/manual.ps.gz>
- Carletta, J. (1996). Assessing agreement on classification tasks: The Kappa Statistic. *Computational Linguistics*, 22(2), 249–254.
- Chi, M., Siler, S. A., & Jeong, H. (2004). Can tutors monitor students' understanding accurately? *Cognition and Instruction* 22: 363–387.
- Chi, M., Siler, S. A., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001). Learning from human tutoring. *Cognitive Science* 25(4): 471–533.
- Cho, B. I. (2000). *Dynamic planning models to support curriculum planning and multiple tutoring policies in intelligent tutoring systems*. Unpublished Ph.D. Dissertation, Computer Science Department, Illinois Institute of Technology, Chicago, IL.
- Cho, B. I., Michael, J. A., Rovick, A. A., & Evens, M. W. (2000). An analysis of multiple tutoring policies. In G. Gauthier, C. Frasson, & K. VanLehn, K. (Eds.) *Proc. ITS 2000*, (pp. 212–221), Montreal, LNCS 1839, Berlin: Springer.
- Di Eugenio, B. (2000). On the usage of Kappa to evaluate agreement on coding tasks. *Proceedings of LREC*, (pp. 441–444), Athens, Greece.
- Evens, M. W., & Michael, J. A. (2006). *One-on-one tutoring by humans and machines*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Fox, B. (1993). *The human tutorial dialogue project: Issues in the design of instructional systems*. Hillsdale, NJ: Lawrence Erlbaum Associates
- Freedman, R. (2000a). Plan-based dialogue management in a physics tutor. *Proceedings of the Sixth Applied Natural Language Processing Conference*. (pp. 52–59). Seattle, WA.
- Freedman, R. (2000b). Using a reactive planner as the basis for a dialogue agent. *Proceedings of FLAIRS 2000*, (pp. 203–208). Orlando, FL.
- Freedman, R. (2001). An approach to increasing programming efficiency in plan-based dialogue systems. In J. D. Moore, C. L. Redfield, and W. L. Johnson (ed.), *Proceedings of Artificial Intelligence in Education*. Amsterdam: IOS Press. 200-209.
- Glass, M., Kim, J. H., Evens, M. W., Michael, J. A., & Rovick, A.A. (1999). Novice vs. expert tutors: A comparison of style. *Proceedings of Midwest Artificial Intelligence and Cognitive Science Conference 99* (pp. 43–49). Bloomington, IN.
- Graesser, A. C., Person, N. K. & Magliano, J. P. (1995). Collaborative dialogue patterns in naturalistic one-to-one tutoring. *Applied Cognitive Psychology*, 9, 495–522.
- Kim, J. H., Freedman, R., Glass, M. S., & Evens, M. W. (2006). Annotation of Tutorial Goals for Natural Language Generation. *Discourse Processes*, 42(1): 37–74.
- Leinhardt, G., & Greeno, J. G. (1986). The cognitive skill of teaching. *Journal of Educational Psychology*, 78, 75–95.
- Li, J., Seu, J. H., Evens, M. W., Michael, J. A., & Rovick, A. A. (1992). Computer dialogue system (cnds): A system for capturing computer-mediated dialogues. *Behavior Research Methods, Instruments, and Computers (Journal of the Psychonomic Society)*, 24(4): 535–540.
- Michael, J. A., Rovick, A. A., Glass, M. S., Zhou, Y., & Evens, M. (2003). Learning from a computer tutor with natural language capabilities. *Interactive Learning Environments*, 11(3), 233–262.
- Pilkington, R. M. (1999). Analyzing educational discourse: The DISCOUNT scheme Technical Report 99/2. Computer Based Learning Unit, University of Leeds, Jan.1999, <http://cbl.leeds.ac.uk/rachel/papers/Discoun99/>
- Rovick, A. A., & Michael, J. A. (1992). The prediction table: A tool for assessing students' knowledge. *American Journal of Physiology*, 263(6), Pt. 3: S33-S36.
- Shah, F., Evens, M. W., Michael, J. A., & Rovick A.A. (2002). Classifying student initiatives and tutor responses in human keyboard-to-keyboard tutoring sessions. *Discourse Processes*, 33(1), 23–52.
- VanLehn, K., Graesser, A. C., Jackson, G. T., Jordan, P., Olney, A., & Rosé, C.P. When are tutorial dialogues more effective than reading? *Cognitive Science*, 35, (2007), 3–62.

Table 3. Frequency counts of specific tutoring strategies and policies in expert tutoring sessions. Each row represents a separate phase in the tutoring process in a particular session. POL denotes the policy used during the phase, and the other columns give the number of times the individual strategies occurred in the phase.

PHASE	POL	DLR	DET	MVF	STA	EXA	DIA	ANA	ALG	PRE	NGR	SUM	LOG	PRB	ROG	Total
K30DR1	V2	2	0	0	0	0	0	0	0	0	0	0	0	0	0	2
K30RR1	V2	0	2	1	0	0	0	2	0	0	0	0	0	0	0	5
K30SS1	V2	0	0	0	0	0	1	1	1	0	0	1	0	0	0	4
K30DR2	V2	0	0	2	0	0	1	0	0	1	0	0	0	0	0	4
K30RR2	V2	0	1	0	0	0	0	2	0	0	2	1	0	0	0	6
K30SS2	V2	0	0	0	0	0	0	0	0	0	1	0	0	0	0	1
K31DR1	V2	1	1	2	0	0	0	2	0	1	0	2	0	0	0	9
K31RR1	IMF	0	0	2	0	1	0	0	0	0	1	0	1	0	1	6
K31DR2	IMF	0	0	5	0	0	1	0	0	0	0	0	0	0	0	6
K32DR1	V2	0	0	2	0	0	1	0	0	0	0	0	1	1	0	5
K32RR1	V2	0	0	0	0	1	2	0	0	0	0	0	0	0	0	3
K32SS1	V2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
K32DR2	V2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
K32RR2	V2	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
K32SS2	V2	0	0	0	0	0	0	1	1	0	0	0	0	0	0	2
K33DR1	V2	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
K33RR1	V2	0	1	0	0	0	1	0	0	1	0	0	0	0	0	3
K33SS1	V2	0	1	0	0	0	0	0	0	0	0	0	0	0	0	1
K33DR2	V2	0	0	3	0	0	0	0	0	0	0	0	0	0	0	3
K33RR2	V2	0	0	0	0	1	0	0	0	0	1	0	0	0	1	3
K33SS2	V2	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
K34DR1	V2	1	0	0	0	0	0	2	0	0	0	1	0	0	0	4
K34RR1	IMF	0	2	5	1	0	0	2	0	2	0	1	2	0	0	15
K34DR2	IMF	1	1	5	0	0	3	0	0	0	0	2	0	0	0	12
K35DR1	V2	0	0	0	0	0	1	0	0	0	0	0	0	0	0	1
K35RR1	V2	0	0	0	0	1	1	0	0	0	0	1	0	0	0	3
K35SS1	IMF	0	0	4	1	0	0	0	1	0	0	0	0	1	0	7
K35DR2	IMF	1	0	6	0	0	1	0	0	1	0	1	0	0	0	10
K35RR2	IMF	0	1	3	0	0	0	0	0	0	1	1	0	0	1	7
K35SS2	IMF	0	0	5	0	0	0	0	1	0	0	0	0	0	0	6
K36DR1	IMF	0	0	6	0	0	1	0	0	0	0	0	0	1	0	8
K36RR1	IMF	0	1	9	1	1	1	0	0	0	1	1	0	0	1	16
K36DR2	IMF	0	0	4	0	0	0	0	0	0	0	0	0	0	0	4
K36RR2	V2	0	1	1	0	0	2	1	0	0	0	0	0	0	0	5
K36SS2	IMF	0	0	6	1	0	2	0	1	0	0	0	0	1	0	11
K37DR1	V2	1	0	0	0	0	1	0	0	0	0	1	1	1	0	5
K37RR1	V2	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
K37SS1	V2	0	0	0	0	0	0	0	0	0	1	1	0	0	0	2
K37DR2	V2	0	0	3	0	0	2	0	0	0	0	1	1	0	0	7
K37RR2	V2	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1
K37SS2	V2	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
K38DR1	IMF	2	0	4	0	0	0	0	0	0	0	1	2	0	0	9
K38RR1	IMF	0	0	6	1	1	1	0	0	0	2	1	1	0	0	13
K38DR2	IMF	0	0	5	0	0	1	0	0	0	0	0	0	0	0	6
K38RR2	V2	0	0	1	0	0	1	0	0	0	1	1	0	0	0	4
K38SS2	V2	0	0	0	0	0	1	0	0	0	1	0	0	0	0	2
Total	V2	5	7	15	0	3	15	11	2	3	7	10	3	6	2	89
Total	IMF	4	5	75	5	3	11	2	3	3	5	8	6	3	3	136
TOTAL	ALL	9	12	90	5	6	26	13	5	6	12	18	9	9	5	225