

Feedback on children's stories via multiple interface agents

Judy Robertson¹ and Peter Wiemer-Hastings^{2*}

¹ University of Edinburgh
ICCS/Division of Informatics

2 Buccleuch Place
Edinburgh EH8 9LW Scotland

² DePaul University
School of Computer Science,
Telecommunications, and Information Systems
243 South Wabash Avenue
Chicago IL 60604 USA

Abstract. This paper describes StoryStation, an intelligent tutoring system designed to give ten to twelve year old children feedback on their creative writing. The feedback is presented via eight animated interface agents. Each agent gives a different sort of support to the writer including: a thesaurus, a dictionary, feedback on vocabulary and characterisation, help with spelling, help with plot structure, example stories to read and help with the interface itself. This paper focuses on the strategies for generating feedback to the children and discusses some issues in presenting this feedback through the interface agents.

1 Introduction

Writing stories is hard work. Many writers find it stressful; they might worry about whether an idea is any good, whether the story will interest the reader, whether their spelling is adequate or even whether their handwriting is legible. Children need a lot of encouragement, help, support and feedback during the writing process. After investing so much effort in writing a story, each pupil deserves to have it read and appreciated. Unfortunately, even the best teachers don't have time to give each pupil the support she needs.

A possible solution to the problem is to use an intelligent tutoring system to give the children help and feedback on their stories. A tutoring system can never be an appreciative, amused, scared or sympathetic audience but it can create more time for the teacher to be an appreciative reader. The tutoring system can give lower level suggestions and support to children while the teacher is engaged in high level discussions about stories with other pupils.

* This work was supported by grants from the James S. McDonald Foundation and from the EPSRC. The pilot study reported here was carried out in conjunction with Charlotte Moss. Thanks are also due to the pupils and staff of Sinclairtown Primary School and St Columba's Primary School, Fife, Scotland.

StoryStation is an intelligent tutoring system designed to provide support and feedback to children on a variety of writing skills, based on a prototype described in (Wiemer-Hastings and Graesser, 2000). The system is outlined in Section 3, and the mechanism for supplying feedback to the pupils is described in more detail in Section 4. Much of the feedback is acknowledgement and praise for good work; through recognition of the children's best efforts at mastering writing techniques, StoryStation encourages them to use the same skills in the future. StoryStation provides dictionary and thesaurus facilities, a bank of example stories written by other children, help with spelling, plot structure, and vocabulary usage and techniques for portraying story characters. Each of these features is presented to the user through an animated interface agent. Figure 1 shows the interface to StoryStation.

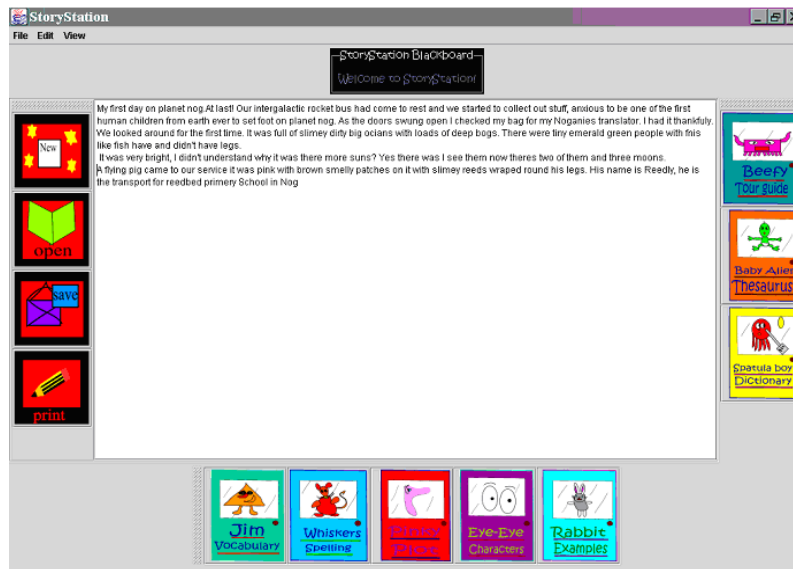


Fig. 1. The interface to StoryStation

Flower (1994) described how important it is for students to learn strategies to manage the multiple constraints involved in writing. She characterized writing as the “negotiated construction of meaning”, and described a variety of voices that speak during the process of composition. Voices correspond to the different types of constraint that impinge upon the process. The term voice emphasises that the constraints are not innocent bystanders to the process. Instead, they are actively involved, pushing the writer in different directions. Thus, the writer must “negotiate” with these voices to achieve a solution which creates meaningful text, and (at least partially) satisfies the constraints. The biggest problem is that the pupil is not normally consciously aware of the voices. This project’s goal is to

help the writer convert these ethereal forces into concrete considerations that she is consciously aware of and then can reason about. Thus, in StoryStation, we have chosen to embody different aspects of writing as unique animated pedagogical agents.

Instead of an “all knowing” computer telling the pupil that something they have done is incorrect, several characters give their “opinions” on the pupil’s composition. Different characters can also disagree with each other. Thus, pupils can learn about the constraints inherent in writing and the interactions between them, and build strategies for managing them.

StoryStation has been developed in a conjunction in close consultation with pupils and teacher at a local state funded primary school. We have used a child centred design strategy in which a team of eight children and two teachers have helped the researchers at all stages of the project. One of our collaborators is a retired teacher with forty years of teaching experience; her expertise and knowledge of the Scottish National Curriculum have been invaluable to the project. We have also worked with primary six and seven classes in three other schools while gathering story corpora and testing aspects of StoryStation on subjects outside our design team. We have chosen to work in state funded schools with pupils of a range of ability levels and socio-economic backgrounds. The software is still under development, but as part of our development methodology we have conducted formative evaluation of some aspects of it. These studies are outlined in Section 5. In this paper we focus on some issues concerning the presentation of feedback to the pupils, and the strategies we have developed for this. We will empirically compare the effectiveness of these strategies in future work.

Agent interfaces may seem intuitively appealing, especially for children’s software, but what effect do such interfaces have on their motivation and learning? As discussed in Section 2, there is currently no conclusive or consistent evidence that they are effective. We intend to explore these questions in future studies using StoryStation. The paper concludes with some lessons learned from the pilot work so far, with some implications for future research in this area.

2 Background

Several computer writing environments were developed in the late 1980s (Britton and Glynn, 1989; Sharples et al., 1988). The design of these environments was informed by research on the cognitive processes during the composition process (e.g. Flower and Hayes, 1980). Such writing environments provided appropriate representations of the composition at different stages of the process (e.g. planning, generating, revising) and the facility to switch between the representations in order to cater for different writing styles.

More recently, research into writing systems has refocused on providing emotional as well as cognitive support for writers. T’riffic Tales (Brna, Cooper and Razmerita, 2001) is a cartoon creation program designed for five year olds who are learning to write. Pupils can ask for help with their stories from an emotional pedagogical agent named Louisa. The design of Louisa was informed by research

into empathic interactions between class teachers and their pupils. The designers take the view that a pedagogical agent should provide emotional support — care, concern continuity, and security — for learners as well as cognitive, domain related support. They propose a cycle of behaviours for such an agent which signals to the pupil that help is available, provides help in an engaging, interactive manner, and ensures the pupil that further help is available if needed (Brna, Cooper and Rasmerita, 2001).

This model seems particularly appropriate in the writing domain. Bereiter and Scardamalia (1982) report empirical evidence on the effects of encouraging and prompting novice writers. It was found that children would write more on a subject if they were initially asked to write as much as possible and were prompted to continue at points when they claimed to have nothing more to write. These very simple interventions tripled writing output. The prompts were content-free motivational prompts such as “You’re doing well. Can you write more?”. Bakunas (1996) found that discourse-related prompts which suggested a structure for ideas were even more helpful than purely motivational prompts in the context of generating ideas for an essay. These empirical results suggest that interactions with animated agents which exhibit encouraging, empathic behaviours will be beneficial for writers.

The preceding discussion of appropriate design for pedagogical agents does not focus on the effectiveness of such agents in comparison to “traditional” graphical user interfaces. However, it cannot be assumed that animated pedagogical agents are more effective than GUIs in terms of either students’ motivation or learning. Indeed, Dehn and van Mulken (2000) review empirical studies of the impact of animated agents on the user’s experience with using the software, her behaviour while using it and her performance on the task with which the agent is intended to assist. The authors report that there are few empirical studies which address these issues, and that the results are inconclusive. They conclude that the literature to date “does not provide evidence to for a so-called persona effect, that is, a general advantage of an interface with an animated agent over one without an animated agent” (Dehn and van Mulken; 2000: p. 17). Furthermore, the methodological validity of some studies is questionable. For example, Lester et. al (1997) concluded that the presence of the animated agent in an intelligent tutoring system improved the learners’ problem solving skills. Dehn and van Mulken point out that these conclusions are suspect because there was no control condition that provided the same advice without an animated agent.

Given the lack of evidence due to a small number of studies, some of which are confounded, Dehn and van Mulken call for further methodologically sound studies in this area. We intend to conduct such studies during the course of the StoryStation project.

3 StoryStation overview

StoryStation is being developed using a child centred design methodology, adapted from (Druin, 1999; Scaife and Rogers, 1999). We have worked closely with a team

of eight (ten to twelve year old) pupils in a state funded primary school during the design process. As part of this process, the pupils have evaluated other pieces of writing software, suggested ideas for the software features, designed and animated interface agents, and created icons for the interface (see Wiemer-Hastings and Robertson, 2001). We have also conducted larger scale requirements analysis with two classes, interviewed teachers, and started exploratory work into agent interactions through a Wizard of Oz pilot study (see Section 5).

The software is designed to be most effective when assisting children with a story re-telling task. This activity is used in classrooms as a way of focusing young writer's attentions on writing skills rather than the imaginative skills of creating a story plot. The writer's task is to recount a story he has heard or read before. As he does not have to devote cognitive resources to thinking of a storyline, he can spend more time writing descriptions of the scenes and characters. A story re-telling task is also easier to support in an intelligent tutoring system because there is more information available about what the user is trying to achieve. However, many StoryStation features can be used for any writing task, including factual writing assignments

StoryStation is a work in progress. At the time of writing the backend language processing features are complete, as are the student modelling and teaching rules. The interface is under development, and is not yet integrated with the backend. StoryStation can also be used in a batch mode to process stories from text files and generate feedback files in the absence of a user.

The features of StoryStation are as follows. Firstly, there are some simple facilities to assist the users as they create their stories, but which do not contain intelligent feedback. Interface help is provided in the form of documentation written by the pupils in the design team. There is also a library of example stories, collected from pupils in a variety of local primary schools after visits by storytellers. The purpose of this feature is to give children ideas for their own stories from reading and critiquing stories written by their peers. The same stories will be the basis of writing tasks using StoryStation. StoryStation also provides an interface to the dictionary and thesaurus features of WordNet (Miller et al, 1993) because the classes we have worked with regularly use such resources for writing exercises.

In addition to these simple features, there are some forms of support which require some natural language processing techniques. We describe these in the following paragraphs.

Pupils may have difficulty in writing coherent plots, even when re-telling stories with which they are familiar. Although they are capable of orally re-telling the plots, the difficulty of the writing task interferes with the re-telling, resulting in incoherent stories. Furthermore, certain writing techniques are best used in particular plot episodes. For example, lengthy character descriptions generally work better at the beginning of a story because descriptions would slow down the pace of exciting events. As another example, a writer might use shorter sentences to make an action sequence more exciting (Corbett, 2001). It is possible to supply help and suggestions on plot structure if the user is writing a

known story. By setting the user a story re-writing task, a reasonable strategy for reducing cognitive overload, StoryStation can use an algorithm based on latent semantic analysis (Landauer and Dumais, 1997) to match the user's plot episodes with episodes in the model story. In this way, it can remind the pupil of plot episodes which she might have forgotten, help the pupil improve coherence by re-ordering episodes, and suggest appropriate techniques for improving particular episodes.

Spelling help was the most frequently requested feature during the requirements gathering exercise. In spite of a move towards the process approach to writing, where a drafting process encourages children to focus on spelling in the later stages of editing, many children are constantly worried by their spelling difficulties. During a recent Wizard of Oz pilot study we observed that lower ability pupils found constant help with spelling reassuring. StoryStation is integrated with Microsoft Word's spell checker. We have found this spell checker to be 80-85% reliable when checking children's stories. StoryStation also supplies spelling support in the form of personalised word banks of commonly misspelled words, specialist vocabulary, and language suitable for particular stories. These have been adapted from vocabulary notebooks which are used by the classes we have worked with for the same purpose.

The vocabulary and characterisation features of StoryStation adopt the strategy of praising the pupils for good use of particular techniques by highlighting "good words" in the pupil's story. "Good words" is a classroom shorthand for long words, unusual words, or words which are not within the pupil's everyday vocabulary. It can also highlight phrases where the pupil has described a character's appearance, personality or feelings, and good dialogue segments. In the classroom, praising a pupil for using a writing technique was until recently less common than reprimanding her for mistakes. However, the schools involved in the project have recently adopted a writing scheme which specifies that teachers should not correct a story in the absence of the author; instead she should highlight parts of the story she enjoyed. Later, in the presence of the author, she can discuss the mistakes and offer suggestions. This approach is considerably more motivating than traditional approaches to marking (see Dunsbee and Ford, 1980) and the class teacher reports that it is working well.

Storystation identifies "good words" by checking the familiarity of each story word in a corpus. The corpus is derived from the mean familiarity and age of acquisition data from the MRC Psycholinguistic database (Coltheart, 1981), and frequency statistics from the British National Corpus (BNC). It can also identify alliteration and minimal pairs (both examples of manipulating language for poetic effect) using the phonetic information in the MRC.

The feedback on characterisation techniques is based on a story characterisation scheme described in Robertson (2000). StoryStation checks the user's story for words which are associated with descriptions of characters' appearance, personalities, feelings and speech. Word lists for each of these categories were derived using WordNet. First of all, seed words for each category were collected from the Linguistic Inquiry and Word Count (LIWC) corpus (Pennebaker et al., 2001);

stories which had been human-rated using the story analysis scheme (Robertson and Good, 2001); and the researcher’s intuition. These word lists were then extended by searching for synonyms, hyponyms and hypernyms in WordNet. The word lists were hand filtered and refined on several story collections. The resulting software was then tested on a fresh corpus which had previously been human-rated using Robertson’s story analysis scheme. Preliminary analysis indicates that inter-rater reliability is acceptable.

StoryStation’s method of generating feedback is described in the next section.

4 Generating feedback

StoryStation generates feedback for the pupils using its student models, curriculum models and some heuristics encapsulated in tutoring rules. It builds a student model for the current user’s story with entries for spelling, word count, number of different words, mean BNC frequency, proportions of nouns, adverbs, adjectives and pronouns, counts of connective usage (such as “and”, “because”, “but”), and counts of characterisation techniques. The tutoring rules use this information to decide what feedback to give on each aspect of the story. Firstly, if the user has previously asked for feedback during the current session, the tutoring rules will compare the current student model to the previous model for the same session to see whether the pupil has improved the story by taking StoryStation’s advice. Secondly, if the user has not asked for help on this story before, the tutoring rules will compare the current student model with the user’s most recent student model from a previous session. Thirdly, if this is the user’s first session with StoryStation, the tutoring rules compare the fields in the user’s current student model to expected values from the curriculum models. The curriculum models contain the norms for each of the five Scottish National Curriculum levels. These were derived by using processing a corpus of 140 stories for each of the linguistic measures stored in the student model and finding the quintiles for each measure. The quintiles provide threshold values for each curriculum level. StoryStation keeps a record of the curriculum level which a child is working towards (this information is supplied by the teacher), and can therefore compare every field in the student model against the threshold value for the appropriate curriculum level.

Note that the third method of comparison is the least reliable because it assumes that the corpus of 140 stories are representative of the children’s stories in general. StoryStation’s performance will become more reliable over time because it will be able to use the user’s previous performance as a benchmark.

The tutoring rules are used to generate overall and specific feedback. The overall feedback on the story gives the user general information about her progress in spelling, vocabulary and characterisation. For example, it could encourage the pupil by mentioning that her spelling has improved since last time, or suggest that she could make her story even better by using the thesaurus to find good vocabulary. These indications of progress will also be useful to the teacher.

StoryStation can also give specific feedback to the pupil at a word or phrase level. It can identify spelling mistakes, offer spelling suggestions, and highlight other features which require the user's attentions. It highlights "good words", and use of characterisation techniques in order to praise the child and encourage her to try using similar techniques in her next story. It can also highlight repetitive use of language, for example overuse of "and", and give general suggestions on how the user could improve this. A discussion of the success of the feedback strategies in pilot work is given in the next section.

5 Pilot results and implications for future work

As part of the design process, we have consulted with pupils and teachers throughout the project. Two informal pilot studies are of interest here - a Wizard of Oz study to explore children's reactions to feedback from animated agents, and a field study to discover how children respond to the StoryStation feedback.

The purpose of the Wizard of Oz study was to discover what pupils thought of the animated agents and whether they took their advice. In this case, the advice was actually typed in by a teacher or another pupil, but appeared to the user as if it came from the agents. Interviews with eight pupils indicated that the pupils mainly appreciated the agent's advice and enjoyed their support. When asked how the agents made them feel, one pupil mentioned: "It made me feel more confident. You know you're not making mistakes in words". Another said "It made me feel happy because it was helping me with my spelling, and the words I didn't know".

One pupil did not enjoy his experience with the agents and seemed upset by one of their comments. The "wizard" who supplied his advice was another pupil, and he bluntly stated that the story did not make sense. When asked about his experiences with the agents, the pupil said "[the advice] made me feel 'Oh, I must be rubbish at my writing then'". It appeared that the pupil had never considered whether it was wise to take the agent seriously or not; he had assumed that the agent had some authority and so was upset by its censure. This implies that careful phrasing of advice is required to avoid demoralising pupils.

A further field study was conducted to see how the children responded to feedback generated by StoryStation. This feedback came in the form of a printed page containing general comments and an annotated version of the pupil's first draft of a story. Twenty three pupils were given this form of feedback on their stories and were asked to redraft the stories after reading the advice.

The pupils did make the changes suggested in the feedback. They worked systematically through the annotated story and changed spelling mistakes and corrected overuse of the connective "and". They appeared to take less notice of the general advice at the start of the feedback and were slightly puzzled by the fact that the purpose of some of the highlighting was simply to praise them. However, once the researchers read out the highlighted praise, reinforcing that the pupils had done well, the pupils understood and were pleased. During

interviews with two groups of five pupils, it seemed as though the more able, more articulate pupils appreciated the feedback more than the less able, shyer pupils. This is possibly because meta-level skills for understanding feedback and responding accordingly are less developed in less able writers.

Some pupils were confused by the colour coding of the comments, but others quickly learned that they only needed to make changes for red and orange highlights, saying “you can just ignore the rest”. One boy who had difficulty with spelling pointed out that the colour coding was demoralising for him. He said “I know I’m not that good at spelling but I like to try to put in good words, but then my whole page gets coloured in yellow”. He was discouraged that his best efforts lead to a lot of negative seeming feedback. This suggests that more spelling support is required during the writing process, and that the feedback strategy for poorer spellers should be modified.

The results of this pilot study suggest that pupils are not accustomed to receiving detailed specific praise on aspects of their writing. To gain the most from StoryStation, pupils may need some help in changing their expectations about story “corrections”. They should be taught to respond to positive as well as negative feedback, and to value their own work.

It may also be the case that more able pupils are better equipped to respond appropriately to advice on their writing, and that less able writers need further support in interpreting and acting on advice. Future studies will explore whether ability is a factor in the effectiveness of advice presented by agents and a traditional graphical user interface.

6 Conclusions

StoryStation is a unique writing environment for children that provides significant support for a story-retelling task. By involving pupils and teachers in every stage of the design process, we have created a system that is well-tailored to the pupils’ preferences and writing processes. StoryStation provides both positive and constructive negative feedback to the pupils. By associating the different types of feedback with different agents, we hope to help the pupils learn the metacognitive skills required to manage the different constraints on the writing process.

In future research, we will use StoryStation as a testbed for addressing a number of important questions regarding how best to support children’s writing. In particular, we can examine the effects of animated pedagogical agents on learning and motivation. Preliminary results have already shown the feasibility of our approach, and suggested new research questions.

References

- Bakunas, B. (1996). Promoting idea production by novice writers through the use of discourse-related prompts. *Applied Psycholinguistics*, 17, pages 385-400.

- Bereiter, C. & Scardamalia, M. (1982). From Conversation to Composition: The Role of Instruction in A Developmental Process. In Glaser, R., Editor, *Advances in Instructional Psychology Volume 2*. Lawrence Erlbaum Associates.
- Britton, B. and Glynn, S. (Eds) (1989). *Computer Writing Environments: Theory, research and design*. Lawrence Erlbaum Associates, New Jersey.
- Coltheart, M. (1981). "The MRC Psycholinguistic Database" *Quarterly Journal of Experimental Psychology*, 33A, 497-505.
- Corbett, P. (2001). *Write your own chillers: create your own spine-tingling stories and terrifying tales*. Belitha Press, London.
- Dehn, D. and van Mulken, S. (2000). "The impact of animated interface agents: a review of empirical research". *International Journal of Human-Computer Studies* 52 1-22.
- Druin, A. (1999) "Cooperative Inquiry: Developing New Technologies for Children with Children". *Proceedings of CHI '99*. ACM Press.
- Dunsbee, T. and Ford, T. (1980). *Mark My Words: A study of teachers as correctors of children's writing*. NATE, London.
- Ellis, S. and Friel, G. (1998). 5-14 Teacher Support: Learning to Write, Writing to Learn: Imaginative Writing. Scottish Consultative Council on the Curriculum.
- Flower, L. (1994) *The construction of negotiated meaning: A social cognitive theory of writing*. Southern Illinois University Press, Carbondale, EL, 1994.
- Flower, L. and Hayes, J. (1980). The dynamics of composing: making plans and juggling constraints. In Gregg, L.W. and Steinberg, E. R., Editors, *Cognitive Processes in Writing*. Lawrence Erlbaum Associates.
- Graves, D. (1983). *Writing: Teachers & Children At Work*. Heinemann Books.
- Hietala, P. and Niemirepo, T. (1998). "The Competence of Learning Companion Agents" *International Journal of Artificial Intelligence in Education* (1998) 9, 178-192.
- Landauer, T.K. and Dumais, S.T. (1997) "A solution to Plato's Problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge" *Psychological Review* 104:211-240.
- Lester, J. Converse, S., Kahler, S, Barlow, S., Stone, B., Bhoga, R. (1997). "The Persona Effect: Affective Impact of Animated Pedagogical Agents" in *Proceedings of CHI 97*, Atlanta.
- Miller, G. , Beckwith, R., Fellbaum, C., Gross, D. and Miller, K. (1993). Introduction to WordNet: An On-line Lexical Database. Online publication, available at <ftp://ftp.cogsci.princeton.edu/pub/wordnet/5papers.pdf>
- Pennebaker, J.W., Francis, M.E., & Booth, R.J. (2001). *Linguistic Inquiry and Word Count (LIWC): LIWC2001*. Mahwah, NJ: Erlbaum Publishers.
- Robertson, J. (2000). The effectiveness of a virtual role-play environment as a story preparation activity. PhD thesis, Edinburgh University. Available at www.cogsci.ed.ac.uk/~judyr .
- Robertson, J. and Good, J. (submitted). "Using a Collaborative Virtual Role-Play Environment to Foster Characterisation in Stories". Submitted to *Journal of Interactive Learning Research*.
- Scaife, M. and Rogers, Y. (1999). "Kids as Informants: Telling us what we didn't know or confirming what we knew already?" in Druin, A. (Ed). *The Design of Children's Technology: How we Design, What we Design and Why*. Morgan Kaufman.
- Wiemer-Hastings, P. and Graesser, A. (2000). "Select-a-kibitzer: A computer tool that gives meaningful feedback on student compositions." *Interactive Learning Environments*, pp 149-169.
- Wiemer-Hastings, P. and Robertson, J. (2001). "Story Station: Children's story writing software". Poster at AIED 2001.