

Generating Stories with Morals

Margaret Sarlej and Malcolm Ryan

School of Computer Science and Engineering
University of New South Wales
Sydney, NSW, Australia
{msarlej,malcolmr}@cse.unsw.edu.au

Abstract. Morals play an important role in why storytelling developed, and help provide stories with structure. We describe a storytelling system which generates short stories that convey one of six common morals identified in Aesop's fables. Morals are represented in terms of patterns of character emotions that arise during the course of a story. To evaluate the system's effectiveness, we compare system-generated stories with human-authored stories and random event sequences. We find system-generated stories convey morals significantly better than random.

Keywords: morals, emotion, storytelling, occ theory.

1 Introduction

Storytelling has long been an integral aspect of human culture, particularly as a mechanism for education and conveying important information [1]. To achieve this, most traditional stories and fables were imbued with a message or lesson: the story's moral. This aspect of storytelling can be leveraged by storytelling systems as a framework for story structure, paving the way for interactive stories which adapt to convey appropriate morals based on readers' choices. Such systems could have a significant impact in education, particularly for children. Here we take the first steps towards this goal, by building a storytelling system that generates stories with morals. To achieve this we need to represent morals. Dyer's Thematic Abstraction Units (TAUs) [2] were developed for concepts similar to morals, but have quite a complex structure. We seek a representation using simpler constructs, and propose the use of sequences of character emotions.

2 Related Work

There has been considerable work in applying emotions to storytelling systems. However, this has mostly been for developing character agents [3–5] rather than the structure or plot of a story. A notable exception is *Mexica* [6], which uses emotional links and tensions between characters for plot construction, but it draws on only two emotions. A relationship between emotions and morals has similarly been proposed in existing work [7, 8], but again the focus was on agent design. In our work we use a broad range of emotions for planning story trajectories, rather than governing agent behaviour.

3 Moral Storytelling System (MOSS)

Even very simple stories require modelling three elements: action, character and plot [9]. The Moral Storytelling System (MOSS) consists of three layers: Action, Emotion and Moral, which correspond to action, character and plot respectively. It was built using Answer Set Programming (ASP) [10], with each layer implemented declaratively as a logic program. We used the Potassco suite of Answer Set solving tools [11] to generate solutions. A Perl script converts the output into a text description of the events, their outcomes and characters' emotions.

3.1 Action Layer

The Action Layer models the physical aspects of the story world. This includes characters, properties of the world (fluents), available actions, their effects, and any restrictions on when actions can be performed, all encoded as ASP rules and constraints. We distinguish between *fluents* (properties of the world that can be true or false) and *consequences* (outcomes of actions, corresponding to changes in value of fluents). To demonstrate that our moral rules can be applied across multiple domains we implemented three distinctly different story worlds.

3.2 Emotion Layer

Of the many emotion theories proposed over the years [12–15] we chose to base our work on the OCC theory [16], which was designed with computational modelling in mind. Table 1 lists the OCC emotions and their definitions in MOSS. Our implementation is largely based on Adam, Herzig and Longin's logical formalisation [17], thus we use their terminology to present our definitions:

- $Bel_i(c)$: Agent i believes consequence c is currently true.
- $Expect_i(c)$: Agent i considers consequence c to be probable and does not believe it is currently true.
- $Des_i(c)$: Agent i considers consequence c to be desirable.
- $Idl_i(a)$: Agent i considers action a to be ideal.

Adam, Herzig and Longin exclude Love and Hate, which are used in the OCC theory to define other emotions; our definitions are more faithful to the original OCC descriptions in this regard. To keep our belief model simple, characters' desires and ideals do not change as a story progresses, and we assume omniscience.

3.3 Moral Layer

Many morals of varying complexity appear in stories. The morals we investigate are derived from Aesop's fables [18], which we categorised by moral in previous work [19]. In this study we deal with 6 morals: Retribution, Greed, Pride, Realistic Expectations, Recklessness and Reward. We developed rules for each moral in terms of the OCC emotions based on an analysis of the relevant fables. Figure 1 shows the rules for each moral. The time-points are not necessarily consecutive, but in general $T1 < T2 < T3$. For agent-based emotions, if no target agent is specified the emotion can be felt towards any agent in the story.

Table 1. The OCC emotions and how they are defined in MOSS

Emotion	Type	MOSS Definition
$Joy_i(c)$	Event-based	$Des_i(c)$ and $Bel_i(c)$
$Distress_i(c)$	Event-based	$Des_i(\neg c)$ and $Bel_i(c)$
$Hope_i(c)$	Event-based	$Des_i(c)$ and $Expect_i(c)$
$Fear_i(c)$	Event-based	$Des_i(\neg c)$ and $Expect_i(c)$
$Satisfaction_i(c)$	Event-based	$Joy_i(c)$ and previously $Expect_i(c)$
$FearsConfirmed_i(c)$	Event-based	$Distress_i(c)$ and previously $Expect_i(c)$
$Disappointment_i(c)$	Event-based	$Distress_i(c)$ and previously $Expect_i(\neg c)$
$Relief_i(c)$	Event-based	$Joy_i(c)$ and previously $Expect_i(\neg c)$
$HappyFor_{i,j}(c)$	Event-based	$Bel_i(c)$ and $Bel_i(Des_j(c))$ and $Love_i(j)$
$Pity_{i,j}(c)$	Event-based	$Bel_i(c)$ and $Bel_i(Des_j(\neg c))$ and $\neg Hate_i(j)$
$Resentment_{i,j}(c)$	Event-based	$Bel_i(c)$ and $Bel_i(Des_j(c))$ and $Hate_i(j)$
$Gloating_{i,j}(c)$	Event-based	$Bel_i(c)$ and $Bel_i(Des_j(\neg c))$ and $Hate_i(j)$
$Admiration_{i,j}(a)$	Agent-based	$Idl_i(a)$ and j successfully performs a
$Reproach_{i,j}(a)$	Agent-based	$Idl_i(\neg a)$ and j performs or attempts a
$Pride_i(a)$	Agent-based	$Idl_i(a)$ and i successfully performs a
$Shame_i(a)$	Agent-based	$Idl_i(\neg a)$ and i performs or attempts a
$Love_i(j)$	Object-based	$Admiration_{i,j}(a)$
$Hate_i(j)$	Object-based	$Reproach_{i,j}(a)$
$Gratification_i(a, c)$	Compound	$Pride_i(a)$ and $Joy_i(c)$
$Remorse_i(a, c)$	Compound	$Shame_i(a)$ and $Distress_i(c)$
$Gratitude_{i,j}(a, c)$	Compound	$Admiration_{i,j}(a)$ and $Joy_i(c)$
$Anger_{i,j}(a, c)$	Compound	$Reproach_{i,j}(a)$ and $Distress_i(c)$

4 Evaluation

We conducted a survey asking participants to read MOSS stories, decide whether each has a moral, and if so select that moral from a list. The proportion of stories classified correctly provides a measure of system performance. The difficulty is story interpretation is extremely subjective. Even human-authored stories would not be expected to score perfectly, but can be considered an upper threshold. Thus our survey incorporates three story types: MOSS stories, human-authored stories, and event sequences without any of the emotion patterns corresponding to MOSS morals (we call these *random*). Authors of the human stories were instructed to work within the MOSS domains, to ensure a consistent scope for all stories, and the MOSS text generation script was used to produce the text, so our results reflect the quality of the story plans rather than the text generation. We randomly selected 210 stories spread across the domains, story types and morals; each participant was shown 9 with the option of responding to more. We expected human-authored stories to perform best, but for MOSS stories to lie closer to human story performance than random event sequences.

We collected 831 story responses from 78 participants, covering each story an average of 4.0 times. In aggregate this yields an average of 138.5 responses per moral. Table 2 presents the data in confusion matrices. Although performance

Time	Retribution Type 1	Retribution Type 2
T1	anger: Agent2 at Agent1	anger: Agent2 at Agent1
T2	distress: Agent1 reproach: Agent1 NOT joy: Agent1 pride: Agent1 disappointment: Agent1 distress: Agent2	anger: Agent1 at Agent2 NOT anger: Agent2

Time	Greed Type 1	Greed Type 2
T1	satisfaction: Agent1 (about X) shame: Agent1 reproach: Agent2 at Agent1 NOT distress: Agent1	satisfaction: Agent1 shame: Agent1
T2	distress: Agent1 (about -X) NOT distress: any other agent	satisfaction: Agent1 shame: Agent1
T3		distress: Agent1 reproach: Agent1 NOT joy: Agent1 distress: any other agent

Time	Pride	Realistic Expectations
T1	pride: Agent joy: Agent (about X)	NOT joy: Agent1 satisfaction: Agent
T2	distress: Agent (about -X) reproach: Agent NOT distress: any other agent	disappointment: Agent NOT satisfaction: Agent

Time	Recklessness Type 1	Recklessness Type 2
T1	distress: Agent satisfaction: Agent NOT admiration: Agent reproach: Agent	fear: Agent
T2		fearsconfirmed: Agent remorse: Agent

Time	Reward Type 1	Reward Type 2
T1	gratitude: Agent2 at Agent1	gratitude: Agent2 at Agent1
T2	gratitude: Agent1 at Agent2 NOT gratitude: Agent2	joy: Agent1 admiration: Agent1 NOT distress: Agent1 shame: Agent1 relief: Agent1 joy: Agent2

Fig. 1. Rules for morals in terms of the OCC emotions

Table 2. Confusion matrices for individual morals

HUMAN-AUTHORED STORIES (WITH MORALS)								
	Retrib	Greed	Pride	Real Exp	Reckless	Reward	NONE	RECALL
Retribution	27	2	0	1	4	1	11	58.7%
Greed	16	19	1	1	2	0	6	42.2%
Pride	12	7	11	0	1	1	14	23.9%
Real Exp	5	0	3	6	4	0	27	13.3%
Recklessness	7	2	0	12	9	0	16	19.6%
Reward	4	0	3	1	0	26	10	59.1%
PRECISION	38.0%	63.3%	61.1%	28.6%	45.0%	92.9%		
MOSS STORIES (WITH MORALS)								
	Retrib	Greed	Pride	Real Exp	Reckless	Reward	NONE	RECALL
Retribution	24	1	1	0	2	2	16	52.2%
Greed	14	12	0	0	0	1	21	25.0%
Pride	11	5	3	3	0	1	21	6.8%
Real Exp	1	3	1	10	1	0	33	20.4%
Recklessness	5	1	0	9	9	3	20	19.1%
Reward	5	2	1	0	0	27	13	56.3%
PRECISION	40.0%	50.0%	50.0%	45.5%	75.0%	79.4%		
RANDOM EVENT SEQUENCES (WITHOUT MORALS)								
	Retrib	Greed	Pride	Real Exp	Reckless	Reward	NONE	RECALL
Retribution	1	2	0	3	1	3	28	2.6%
Greed	0	1	1	0	2	3	33	2.5%
Pride	1	1	3	2	2	6	23	7.9%
Real Exp	1	4	0	3	1	0	26	8.6%
Recklessness	0	1	1	2	1	1	33	2.6%
Reward	1	1	0	1	1	2	32	5.3%
No Filter	6	2	3	3	1	3	31	
PRECISION	10.0%	8.3%	37.5%	21.4%	11.1%	11.1%		

varies between morals, in most cases the human stories had the best recall and precision, and random stories the worst. As expected, MOSS stories generally lie in between, with the exception of Pride (where they perform worse than random) and Realistic Expectations (where they perform better than human).

5 Conclusion and Future Work

The effectiveness of MOSS stories in conveying morals was comparable to human-authored stories in the given domains, particularly relative to random event sequences. The relatively poor performance of many human stories highlights both the restricted nature of the domains and the ambiguity inherent in storytelling. Although further rule refinement is required, we found emotion patterns corresponding to morals were useful for imparting structure to stories. However, there are several areas for future work. Our emotion model encompasses all the OCC

emotions, but not the associated intensity variables. Assigning a numeric valence to emotions would allow more fine-grained control over a story. Data from other collections of fables could be used to improve the moral rules. Permitting incorrect beliefs would let us investigate a broader range of morals, and with richer story domains we could explore the impact of the domain restrictions.

References

1. Ryan, K.: The Narrative and the Moral. *The Clearing House* 64(5), 316–319 (1991)
2. Dyer, M.G.: In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension. The MIT Press, Cambridge (1983)
3. Elliott, C.: The Affective Reasoner: A process model of emotions in a multi-agent system. PhD thesis, Northwestern University (1992)
4. Theune, M., Rensen, S., op den Akker, R., Heylen, D., Nijholt, A.: Emotional Characters for Automatic Plot Creation. In: Göbel, S., Spierling, U., Hoffmann, A., Iurgel, I., Schneider, O., Dechau, J., Feix, A. (eds.) *TIDSE 2004*. LNCS, vol. 3105, pp. 95–100. Springer, Heidelberg (2004)
5. Marsella, S., Gratch, J.: EMA: A process model of appraisal dynamics. *Cognitive Systems Research* 10, 70–90 (2009)
6. Pérez y Pérez, R.: Employing emotions to drive plot generation in a computer-based storyteller. *Cognitive Systems Research* 8(2), 89–109 (2007)
7. Shaheed, J., Cunningham, J.: Agents making moral decisions. In: *ECAI 2008 Workshop on Artificial Intelligence in Games* (2008)
8. Battaglino, C., Damiano, R.: Emotional Appraisal of Moral Dilemma in Characters. In: Oyarzun, D., Peinado, F., Young, R.M., Elizalde, A., Méndez, G. (eds.) *ICIDS 2012*. LNCS, vol. 7648, pp. 150–161. Springer, Heidelberg (2012)
9. Ryan, M., Hannah, N., Lobb, J.: The Tale of Peter Rabbit: A Case-study in Story-sense Reasoning. In: *4th Australasian Conference on Interactive Entertainment, Melbourne, Australia*. RMIT University (2007)
10. Lifschitz, V.: Answer set programming and plan generation. *Artificial Intelligence* 138(1-2), 39–54 (2002)
11. Gebser, M., Kaminski, R., Kaufmann, B., Ostrowski, M., Schaub, T., Schneider, M.: *Potassco: The Potsdam Answer Set Solving Collection*. *AI Communications* 24(2), 105–124 (2011)
12. Izard, C.E.: *Human Emotions*. Plenum Press, New York (1977)
13. Plutchik, R.: *Emotion: A Psychoevolutionary Synthesis*. Harper & Row, New York (1980)
14. Russell, J.A.: A Circumplex Model of Affect. *Journal of Personality and Social Psychology* 39(6), 1161–1178 (1980)
15. Frijda, N.H.: *The Emotions*. Cambridge University Press, New York (1986)
16. Ortony, A., Clore, G.L., Collins, A.: *The Cognitive Structure of Emotions*. Cambridge University Press, Cambridge (1988)
17. Adam, C., Herzig, A., Longin, D.: A logical formalization of the OCC theory of emotions. *Synthese* 168(2), 201–248 (2009)
18. Aesop: *Aesop: The Complete Fables*. Penguin Classics. Penguin Books, London (1998), Translated by Temple, O., Temple, R.; introduction by Temple, R.
19. Sarlej, M., Ryan, M.: Representing Morals in Terms of Emotion. In: *8th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, Palo Alto, California*, pp. 69–74. AAAI Press (2012)