

# Narrative and Ethics\*

Greg Lessard<sup>1</sup> and Michael Levison<sup>2</sup>

- 1 French Studies, Queen's University  
Kingston, Ontario, Canada  
greg.lessard@queensu.ca
- 2 School of Computing, Queen's University  
Kingston, Ontario, Canada  
levison@cs.queensu.ca

---

## Abstract

Most computational narrative models descend from the Proppian tradition, in which characters have fixed traits, including their goodness or badness. However, some literary narratives allow for characters to grow and change. Beyond this, it is also important to take account of the attribution of moral traits and evaluations by readers of narrative. How this occurs is a complex matter. We propose here a model and an environment to assist in its study. Specifically, starting from threaded Directed Acyclic Graphs, we generate short narrative texts in English. These are presented in a web-based environment which allows a reader to express some degree of positive or negative moral evaluation. This is captured and used to drive some fundamental moral decision at the end of the text. We show how this model makes it possible to study in some detail how individual readers respond to texts, how the presence or absence of specific textual structures affects these judgments, as well as the interaction among them, and how the linearity inherent in narrative plays a role in affecting moral judgments. Applications of the model include the empirical analysis of one aspect of reading, as well as the opportunity to study in more detail ethical perspectives themselves, both at the individual and collective level. Finally, by capturing and storing reader reactions, the model provides a means of dynamically planning texts to achieve maximum effect with respect to reader judgments.

**1998 ACM Subject Classification** I.2.4 Representation languages

**Keywords and phrases** Directed acyclic graphs, narrative, ethics

**Digital Object Identifier** 10.4230/OASICS.xxx.yyy.p

## 1 Background

It would be difficult to find two broader fields than narrative and ethics. Each has given rise to an enormous range and depth of previous research, in literary studies, philosophy, ethnology, and so on. To a lesser extent, this is also true of the intersection between the two fields. Boyd [2] argues that narratives provide an evolutionarily advantageous means for humans to test theories of action and evaluation in a constructed context, that of the literary work. Booth [1] argues for the importance of an ethical reading of literary texts, and more specifically that exposure to literary texts has the potential to influence the reader's character, as well as the ethical facets of the act of reading itself.<sup>1</sup> Phelan [6] makes the case that "individual narratives explicitly or more often implicitly establish their own ethical standards in order to guide their audiences to particular ethical judgments".

---

\* This work was funded in part by the Social Sciences and Humanities Research Council of Canada.

<sup>1</sup> Booth attributes to Anthony Burgess the affirmation that "the novel is ... a form steeped in morality".



© Greg Lessard and Michael Levison;  
licensed under Creative Commons License BY

Conference/workshop/symposium title on which this volume is based on.

Editors: Billy Editor, Bill Editors; pp. 1–16

OpenAccess Series in Informatics

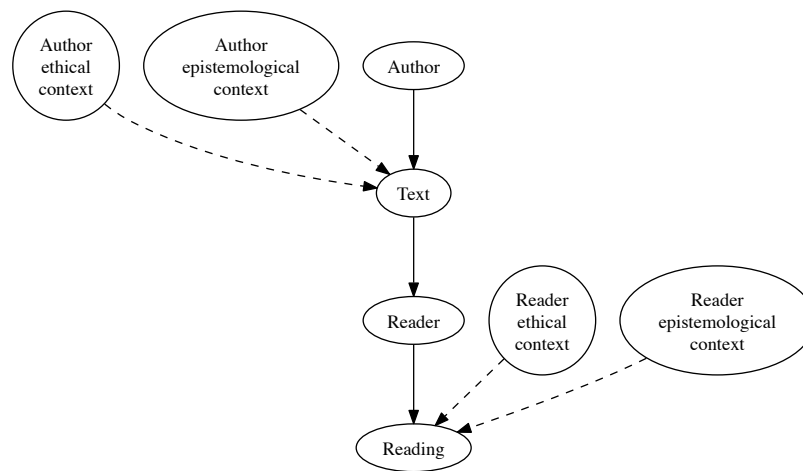


OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

One finds research devoted to the role of ethics in both the production of literary texts (for example, REFs) and in their interpretation (for example, REFs). Narratives have been used as a tool to study the ethical perspectives and issues of various groups. Thus, Gilligan [3], devotes a chapter to the narratives of women faced with issues around abortion. Stadler [7] examines the role of ethics in film.

## 2 Authors, Texts, Readers and Readings

The most basic view of the ‘chain of production’ for narratives involves four elements: an author, a text, one or more readers, and one or more readings made by each reader. This may be represented schematically by the solid lines in Figure 1.



■ **Figure 1** Elements of the narrative chain

In fact, this model leaves out several crucial elements. Literary texts are significantly underspecified;<sup>2</sup> they may be seen as the tip of a semiotic iceberg. Under the water, so to speak, lie at least two unwritten or unspoken dimensions, present for both the author and the reader, represented by dashed lines in Figure 1:

- (a) the epistemological framework – that is, the physical and factual characteristics of the world being described. So, for example, in science fiction, the world presented may have different gravity, there may be different astronomical entities (think of *Star Wars* and its two moons over Luke Skywalker’s home planet). Not all of these are necessarily spelled out explicitly or in detail in the text.<sup>3</sup>
- (b) the ethical framework – that is, the set of cultural and ethical values which are presumed to underlie and to explain or to motivate the actions of characters. As we will see below, these are even more often left implicit.

<sup>2</sup> See, for example, Pavel [5], for discussion.

<sup>3</sup> See, for example, Ryan [?] for an overview of the spatial context of narratives and narrative worlds.

In the simplest cases, an author and his or her readers share the same ethical and epistemological contexts. However, there can also be more or less substantial disconnects. For example, a modern reader of Dante's *Inferno* may not have access to the theology, the social framework and the ethics behind the original. In the extreme case, Dante's description of the various levels of Hell and the inhabitants thereof may take on, for some modern readers, the sense of an arcane travelogue, rather than the exploration of medieval morality, literature, and the politics of Italy at the time, that Dante appears to have intended. To take a case closer to home, the modern reader of Jane Austen's fiction may understand the love story behind one of her novels, but fail to appreciate the financial consequences of spinsterhood that were true in Austen's time.<sup>4</sup>

In what follows, we will be concerned primarily with the interplay between the text, the reader, and the reader's ethical context.<sup>5</sup> As we will see, this relationship may be explored from a variety of perspectives.

### 3 A first case: Androcles and the Lion

For those not familiar with this fable, it tells of an escaped slave, Androcles, who comes upon an injured lion and heals his wounds. Later, the lion and Androcles are captured, separately, by soldiers of the emperor and placed in an arena. However, instead of devouring Androcles, the lion fawns upon him. Puzzled, the emperor asks how this has happened and on hearing the full story, releases both.<sup>6</sup>

It is possible to represent the basic information behind this narrative by means of a *Directed Acycle Graph* (henceforth DAG) composed of nodes and directed edges, which is *threaded* to show the order of narration. This model corresponds to the traditional distinction between *sujzhet* (the set of information which underlies a narrative) and the *fabula* (the set of items presented in some recounting of the narrative). Here, we will refer to these as the *story* and the *narrative* respectively, using the term *narration* to designate the action of threading itself. Since the story (but not necessarily the threading) is a directed acyclic graph, two characteristics are central: (a) the nodes are linked by directed edges, where the subsequent node is seen as depending on the semantic information present on its antecedent node(s); (b) no cycles are permitted in the DAG, although they are possible in the threading. In addition, we will represent the contents of each node by means of *semantic expressions* expressed in a Haskell-like notation.<sup>7</sup>

Using this framework, the basic gist of the Androcles story may be represented by the nodes in lower case in Figure 2 and the edges which link them.<sup>8</sup> Examination of Figure 2 shows that, at at least two points in the story, a character is faced with an ethical choice:

- (a) Androcles could kill the injured lion, or flee him, but instead he heals him;

<sup>4</sup> Of course, one of the functions of literature courses is to elucidate these dimensions and to allow the modern reader to approach a fuller understanding of the original text. Such an approach is necessarily asymptotic.

<sup>5</sup> We will leave aside discussion of the epistemological context for now, although it should be clear that there exists a rich interplay between ethics and epistemology.

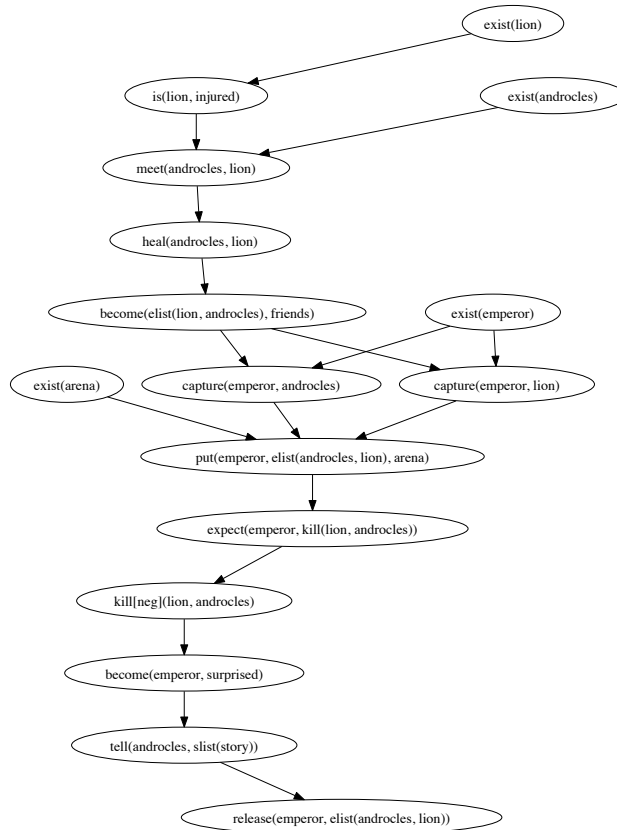
<sup>6</sup> Aesop's *Fables* exist in many versions and many languages. For ease of access to the reader, we use here the version found in Project Gutenberg at <http://archive.org/stream/aesopsfables00028gut/aesopa10.txt>

<sup>7</sup> For details of this notation, see Levison et al. [4].

<sup>8</sup> The DAGS and threading shown here were generated by Graphviz (<http://www.graphviz.org>) from semantic expressions and links between them.

- (b) in the arena, the now-healthy lion could follow his carnivorous tendencies and eat Androcles, but instead he does not, presumably in gratitude for Androcles' earlier unselfish act.

At these points, some unstated ethical principle comes into play and overrides the typical or default behaviour. We can represent one instance of these unstated ethical principles by means of semantic expressions in uppercase, as shown in Figure 3. Here, the principle that friends do not kill friends leads to the lion refraining from killing Androcles.<sup>9</sup>

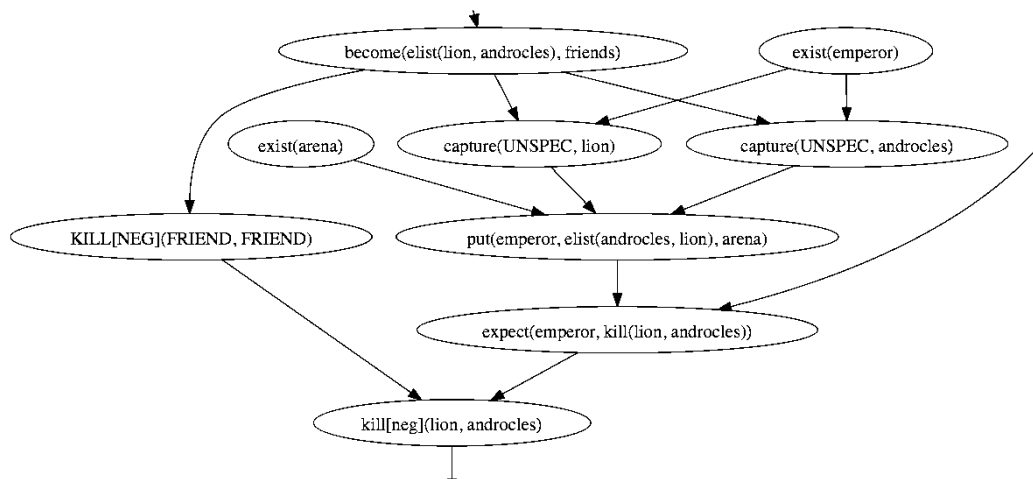


■ **Figure 2** The gist of the Androcles fable shown as a DAG

The fable serves its purpose by presenting a narrative which illustrates the place of these ethical principles in human (and anthropomorphized lion) behaviour. In fact, its role is double. On the one hand, since the reader and the author presumably share the same ethical principles here, the fable ‘makes sense’, and on the other, in its telling, the fable instantiates and emphasizes the importance of these ethical principles.

But what of the case where the reader and the author do not share precisely the same

<sup>9</sup> For simplicity, we will not include here a number of epistemological underpinnings such as the lion’s carnivorous tendencies, his possible hunger, and so on, concentrating rather on the ethical issues. Of course, a fuller model would need to deal with the interaction among these factors, as we will see below. We also include a number of paths which reflect chronological order.



■ **Figure 3** The Androcles DAG with ethical principles included

ethical principles? Also, what is the ‘moral calculus’ which leads to some principle overriding another when two or more are in competition. It is to these questions that we now turn.

#### 4 A second case: the Ant and the Grasshopper

In this fable,<sup>10</sup> an ant and a grasshopper adopt two different strategies during the summertime. The ant works hard to collect food, while the grasshopper plays all day. When winter arrives, the ant has a store of food, while the grasshopper is starving. Although it is not specified explicitly in Aesop’s text, except in part through the moral expressed at the end (One should prepare for hard days ahead), the fable hides an underlying ethical choice facing the ant. This can be represented as the conflict between two principles:

- (1) one should help those in need
- (2) one should take responsibility for one’s own welfare

Aesop’s version of the fable suggests, although this is not stated explicitly, that more weight should be given to principle (2). It is possible to express this relative weighting of the two principles, as well as a test based on these weights, by means of an algorithm in pseudo-code, as shown in Listing 1.

In the fable as written by Aesop, one must assume that  $x$  is less than  $y$  in order to understand the outcome as written.<sup>11</sup> This approach is *interpretive*, in that it takes the textual outcome as primary and attempts to formulate a set of values which justify it. However, it is also possible to imagine a *generative* perspective, in which the reader of a text must provide at the outset some initial weights to the values of  $x$  and  $y$ , thereby prioritizing one of the two principles. We have written a simple program, in awk<sup>12</sup> which asks a user to weight a set of ethical principles on a scale of 0 to 9, where 0 means that the principle

<sup>10</sup> As in the case of Androcles, we have used the Gutenberg version here. However, we will add elements below to enable further ethical analysis.

<sup>11</sup> Note that we have assumed that  $b$  is the default, so that if the values of  $x$  and  $y$  are equal, option  $d$  is chosen. We might equally have provided a middle ground in which no decision can be taken when both values are equal. We will not explore that issue further here.

<sup>12</sup> For details, see <http://www.gnu.org/software/gawk/manual/gawk.html>.

■ **Listing 1** The ant's dilemma

```
# ethical principles
a = "help[should](rich, poor)"
b = "help[should](all, self)"

# available options
c = "feed(ant, grasshopper)"
d = "feed[not](ant, grasshopper)"

# set the parameters
weight(a) = x
weight(b) = y

# test values and choose option
if (x > y) then c, else d
```

■ **Listing 2** Specifications of some ethical weights

```
# ethical principles
"help_strong_weak"|1
"kill_strong_weak"|6
```

is never applied, 9 that it is always applied, which generates the steps of the narrative. A portion of this input is shown in Listing 2, where each line has two fields separated by a vertical bar: first the ethical principle, and second the weight assigned to it.

Once this file is read into an interpreter, the story is run and the values assigned to the principles determine the outcome. A sample run is shown in Listing 3. Once the various ethical principles have been read in from the ethics.dat file, including their various scores, the program SAYS those lines which require no ethical evaluation. When an ethical decision point is reached involving two or more principles, a test is done to determine which of the possible ethical principles has been assigned the highest score. That principle then drives the choice between various courses of action, and the narrative resumes. So, in Listing 3, the fact that the principle of the strong killing the weak outweighs that of the strong helping the weak leads to Androcles killing the lion.

This approach provides the elements of what might be called an *Ethics Workbench*, in which weightings assigned to different principles determine the outcome of some narrative. To the extent that the user of the workbench is content with the outcome generated, he or she will have confirmation of his or her weightings. On the other hand, if one is unhappy with the generated outcome, then the weightings might be revised.

## 5 Further complexities

While the approach just described provides an interesting proof of concept, it has a number of clear weaknesses. First, as described, the model represents only a simple binary choice. Second, it ignores the linearity of narrative, that is, the fact that textual items are not presented all at once. And third, it provides only a single dimension for the ethical decision. In fact, one can imagine various factors which might influence a decision like the one shown above, including:

- (a) the *consequences* of the decision. In the Aesop fable, although this is unstated, the

■ **Listing 3** A sample run of a simple ethics decision program

```

Ethics> awk -f ethics.awk androcles.dat
Reading ethics.dat file...
ethics.dat line no = 1 Principle = tell_truth_always
Score = 9
ethics.dat line no = 2 Principle = hide_painful_information
Score = 1
ethics.dat line no = 3 Principle = kill_strong_weak
Score = 6
ethics.dat line no = 4 Principle = protect_children
Score = 8
ethics.dat line no = 5 Principle = protect_adults
Score = 5
ethics.dat line no = 6 Principle = help_strong_weak
Score = 1
ethics.dat line no = 7 Principle = assume_responsibility_for_self
Score = 2
ethics.dat line no = 8 Principle = show_gratitude_for_service
Score = 8
Done reading ethics file...
Size of ethics.dat 8 lines
=====
Begin analysis of narrative...
SAY exist_androcles
SAY exist_lion
SAY is_lion_injured
SAY meet_androcles_lion
-----
Testing line of narrative file = 6
Value of current line = kill_strong_weak
Proposed action = kill_androcles_lion
Matched line 3 in ethics.dat. Ethics principle = kill_strong_weak
Ethics score = 6
Maxscore now = 6
Choice now = kill_androcles_lion
-----
Testing line of narrative file = 7
Value of current line = help_strong_weak
Proposed action = help_androcles_lion
Matched line 7 in ethics.dat. Ethics principle = help_strong_weak
Ethics score = 1
Score = 1 less than current Maxscore of 6
Choice now = kill_androcles_lion
-----
Lastline now = kill_androcles_lion
SAY kill_androcles_lion
So follow up with... die_lion
SAY die_lion
Ethics>

```

- grasshopper will presumably die of hunger. However, in a variant of this example, if the request is only for a single meal and other food will be available the next day whatever the ant's decision, then the consequences of a night's hunger are perhaps less serious;
- (b) the *players* involved. Although the ant and the grasshopper are anthropomorphized in Aesop's text, they are still insects. At one level, the loss of an insect might be seen as less serious than the loss of, for example, a human being. If the players were humans, would the ant's decision be the same?
  - (c) the *circumstances* surrounding the ant's decision. In the Aesop fable, the grasshopper persists in unproductive behaviour despite the ant's repeated warnings, and in some versions even mocks the ant for its hard work. But what if the grasshopper had only one choice to make and made a bad one? Or what if the grasshopper was simply foolish? Or what if the ant had more than enough food, as opposed to the situation where any donation it made might lead to hunger or starvation for itself or its family?
  - (c) any *mitigating factors*. In the fable, it is assumed that both the ant and the grasshopper are healthy and able-bodied adults. However, what if the grasshopper is old, or ill, or very young? Would this influence the outcome?

In addition to these factors, it is also necessary to deal with the evolution of character traits.

Much of the research on computer based narrative has adopted a perspective based on the Proppian framework (REF), in which characters are treated mostly as static entities which do not change over the course of the narrative. So, for example, the wicked witch remains wicked and does not undergo a transformation. There are some exceptions such as REF, where the narrative is driven by the seeking of goals by one or more characters. However, even in such cases, the character who achieves a goal typically does not change as a result of this.

This unchangingness of characters in computer generated texts actually runs counter to what we find in the majority of human-produced narratives, where character development is the norm (REF Phelan). To take just one example, the character Phil in the movie *Groundhog Day* (REF) begins as self-centred and egotistical, but in the course of the narrative, learns to become thoughtful and generous.<sup>13</sup>

*Groundhog Day* is somewhat special, in that the character Phil finds himself in a temporal loop which causes him to relive the same day (February 2nd), but with the possibility of trying out different strategies for the day. Over the course of many iterations, he evolves from egotism, to despair, and finally to altruism and love, at which point the cycle is broken and he moves on to February 3rd. In many respects, the structure of the film is similar to that of a video game, with an entry point, a quest, and a goal which leads to release.

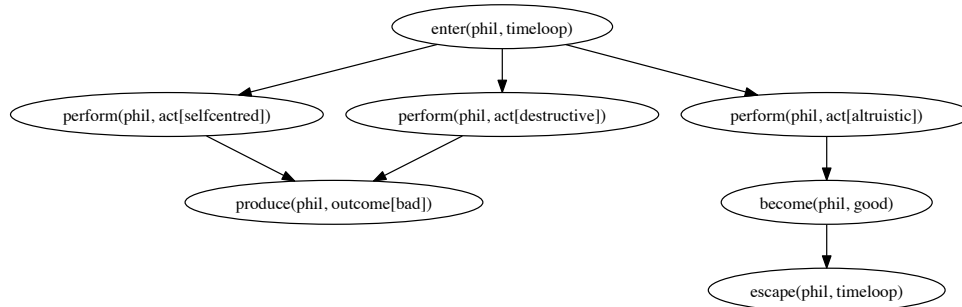
Phil's ethical development may be represented at the most abstract level as a series of trials using the same framework, but adjusting some of the parameters. In one case, he meets a beggar on the street, first ignores him, then gives him money, and finally helps and feeds him. In another, he meets an annoying insurance agent, first escapes from him, then hits him, and finally buys insurance. Or again, on meeting Rita his producer, he first irritates her by his flippancy and self-centred behaviour, then tries to seduce her, and finally falls in love with her.

All of these trajectories can be represented, at a basic level, by the DAG shown in Figure 4.

---

<sup>13</sup> For a detailed analysis of this evolution, see REF.





■ **Figure 4** A high-level DAG for Phil's moral evolution

However, while it illustrates the fact that only one of the possible kinds of behaviour leads to a good outcome, this high-level model fails to capture several important points:

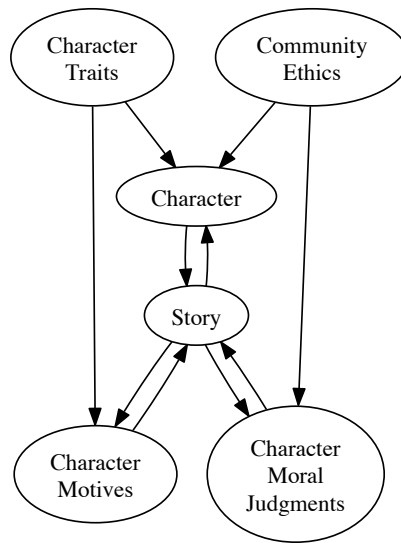
- (a) it ignores the complexities of threading which underlie the narrative. In fact, Phil tries these various approaches with respect to various issues, sometimes numerous times. We will return later to this issue of the interplay between the linearity of a text, which we capture by means of threading, and the architecture of a text, which is captured by a more or less abstract DAG.
- (b) it ignores the fundamental distinction between *acts* and *habits*, which dates back at least as far as Aristotle (REF Nichomachean Ethics). In brief, a single act can be judged to be good or bad, but a person's character or personality can also be judged good or bad. These two can co-exist, in that a good character can perform a bad act, and vice-versa. Equally importantly, according to Aristotle, the repeated performance of good acts leads to habituation and thereby a good character. In that respect, *Groundhog Day* is fundamentally aristotelian. Over the course of a narrative, these two dimensions of acts and habits can be presented in a variety of ways. A character can be presented explicitly and from the outset as good or evil (for example, "Once upon a time there was a wicked witch..."), or alternatively, the reader can come to the conclusion that a character is habitually good or evil based on some recounted series of actions over the course of the narrative. We will return to this below.

## 6 A multi-dimensional model

On the basis of the preceding discussion, we can modify the model ethical factors within a text, as shown in Figure 5.

This model assumes a preliminary distinction between characters and stories (that is, the DAG). This corresponds to the distinction in playwriting between the *Dramatis personae* and the acts of the play.

We assume further that each character may have some traits assigned at the outset by the narrator. This is the case of a character presented as "the wicked witch" for example. Alternatively, a character's traits may be revealed by their actions over the course of the narrative, by their actions, or by comments made by the narrator. As we noted above, in some narratives, the character's traits will remain unchanged (as in fairy tales), while in



■ **Figure 5** An enriched representation of ethical considerations

others they may evolve, as in *Groundhog Day*. In the latter case, we might add a further line from the character node to the character traits node.

We also consider that a character's traits may be evaluated with respect to some set of community ethics. For the moment, we assume that these are the ethics which underlie the world and the situations described in the story. So, the fact that a character is good at killing may be evaluated differently in a narrative world at war, or in the spy genre, than in a narrative world set in a dinner party.

We assume that the various characters bring to the story a set of traits which help to drive the logic of the narrative. In particular, we assume that through the story, character's will participate in events and acquire motives. Thus, in detective fiction, the fact that one character has been swindled by another provides a motive for the first character to do some act of violence on the second. As noted by the double arrows, the story may help create motives, but motives may also drive the logic of the story. In addition, character traits may influence directly character motives, no matter what happens in the story. This is the case of the so-called 'fatal flaw' found in some classical characters.

In a somewhat similar fashion, a character may have moral judgments which apply to some aspect of the story, and these moral judgments may be expressed in the story or lead to some act in the story. Similarly, a character's moral judgments may be bifurcated between what he or she believes to be moral and what the community believes to be so. This is the case of the hero who goes against some community value to achieve what he or she believes to be right.

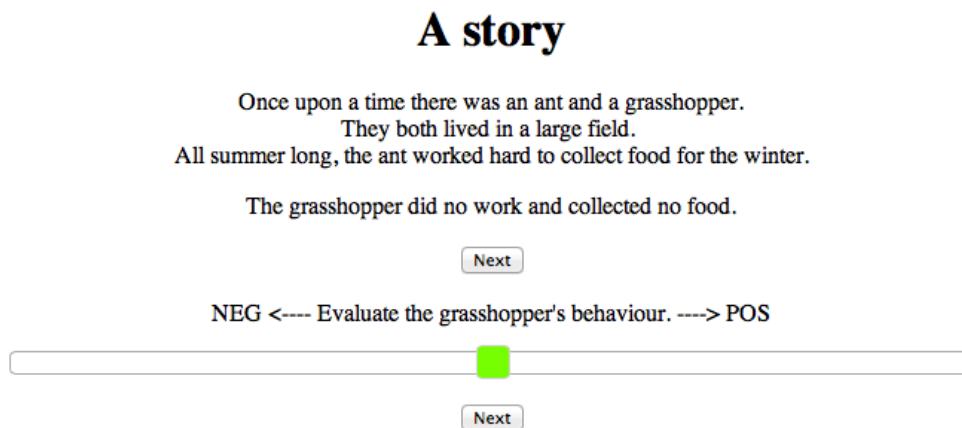
The model we have just presented is in some sense 'narrative internal'. A fuller model

requires that we add to the equation both the author/narrator<sup>14</sup> and the reader<sup>15</sup> In a sense, both author and reader sit ‘outside’ the world we have just illustrated in Figure 5. In principle, the author presumably has access to all items, while the reader’s knowledge is mediated by access to the story, as well as some degree of (possibly limited or inexact) access to the community ethics from background knowledge.

## 6.1 An Ethics Workbench

Using the model described above, it is possible to build a richer framework for testing reader reactions to a text. To do this, we constructed a web page which presents a story in order, one element at a time, and at various points in the story, asks the reader to respond to a certain number of ethical and other questions about the evolving story before moving on to the next story element. Reader judgments were expressed by moving a slider, where movement to the left indicated a progressively greater negative judgment, while movement to the right indicated a progressively more positive ethical judgment.<sup>16</sup> Each slider represents a linear scale, with 0 at the midpoint and -10 and +10 as endpoints. The slider is set initially at the midpoint of the scale.

Once a reader judgment has been entered, the slider disappears, thereby preventing explicit comparison with previous judgments. Figure 6 illustrates the state of the text at a first judgment point:



■ **Figure 6** A screenshot of the Ethics Workbench

Two test scenarios were developed, differing at only one point, as shown below in item (c), where the reason for the grasshopper’s lack of work are presented either as age and illness, or as a desire to sing and mock the ant. The framework for both scenarios is shown below, where sections beginning with letters indicate items of the story as told, and sections beginning with digits indicate requests for judgment on the part of the reader.

<sup>14</sup>For the moment, for the purposes of this illustration, we will conflate the author and the narrator. We recognize that a richer model would distinguish the two. See REF

<sup>15</sup>We also conflate here the *narrated* (REF) and the actual human reader. Again, a richer model could establish differences between them that would take on some importance in at least some texts.

<sup>16</sup>The use of scales to capture subject judgments has been shown to provide rich and reliable data in linguistics (see REF). We have therefore adopted the model here.

- (a) Once upon a time there was an ant and a grasshopper. They both lived in a large field. All summer long, the ant worked hard to collect food for the winter.
- (b) The grasshopper did no work and collected no food.
- (1) *Evaluate the grasshopper's behaviour.*
- (c) Either: He was old and ill. Or: He sang in the sun and mocked the ant for its laborious behaviour.
- (2) *In light of this new information, evaluate the grasshopper's behaviour.*
- (d) At the end of the summer, the ant had a lot of food stored away, but the grasshopper had none. When the first frost came, the grasshopper approached the ant and asked for food.
- (e) The ant told the grasshopper that since he hadn't worked, he deserved no food.
- (3) *How much do you agree with this decision?*
- (f) The grasshopper slowly starved to death.
- (4) *How appropriate is this outcome?*
- (5) *In light of the whole story, evaluate the ant's behaviour.*

## 7 Some preliminary results

The scenarios were presented to 10 adult subjects, 5 male and 5 female. All subjects were given as much time as required to read and to enter their judgments. Half the subjects were exposed to the *old and ill* scenario, and half to the *mocking* scenario. We were interested in particular in four issues:


- (a) Narrative is linear, and subsequent items in a test are invisible to a reader until they are reached. Given this, to what extent can subsequent items 'inflect' previous judgments? Questions 1 and 2 tested this, by asking first for a judgment of the grasshopper's lack of work, without any background information, and then asking again in light of further information on the grasshopper's behaviour.
- (b) Ethical decisions are presumably related, at least to some extent, to presented states of affairs. To the extent that questions 1 and 2 present the 'players' and their respective behaviours, it is possible to ask how much the ant's decision not to share food with the grasshopper can be related to the situation in which the two find themselves, in particular when comparing the two scenarios. Question 3 tests this point.
- (c) Ethical decisions have consequences. In the Aesop's Fable upon which our experiment is based, the typical closing point is the ant's pronouncement that no food will be shared. What this means for the grasshopper is unstated. In the testing here, we have drawn the conclusion that the grasshopper dies. This allows us to test the reader's reaction to this outcome, as is done in question 4.
- (d) More specifically, it is possible to presume that the ant's decision has led to this (although perhaps in conjunction with the grasshopper's behaviour during the summer). As a result, rather than simply representing some alternative behaviour to the grasshopper's, the ant is now an ethical agent, whose behaviour may be evaluated. Question 5 tests the link between the outcome of the ant's decision and the ant's behaviour from the reader's perspective by asking subjects to pronounce on the ant's behaviour.

Figures 7 and 8 show the results of the tests.<sup>17</sup> The left hand axis represents the slider value at each point, the various Series represent the trajectory of judgments for each subject,


---

<sup>17</sup> Since this is preliminary work and the number of subjects is small, no statistical tests were done on the data presented here. The focus is rather on broad tendencies found in the data, as well as an indication of whether the environment used can work in practice.

and the digits from 1 to 5 across the middle of the graph represent each of the five decision points.



■ Figure 7 Subject testing results for *old and ill* scenario



■ Figure 8 Subject testing results for *mocking* scenario

### 7.1 Narrative order and inflection of ethical judgments

Examination of Figure 7 shows that in the *old and ill* scenario, the subsequent explanation of the grasshopper’s inactivity causes all judgments to be inflected in a positive direction, as can be seen by comparing points 1 and 2 on the scale. On the other hand, as Figure 8 shows, 4 of 5 subjects inflect their evaluation of the grasshopper in a negative direction when learning of his singing in the sun and mocking the ant. In the case of Series 2, the subject tested noted that their evaluation was based on the grasshopper’s singing, which was seen as a good thing.<sup>18</sup>

Taken together, this data provides at least *prima facie* evidence that reader judgments follow the ordered presentation of information in a narrative.

### 7.2 Ethical situations and ethical decisions

As noted above, the two scenarios provide background and then ask the reader to pronounce on his or her agreement with the ant’s decision. Analysis of Figures 7 and 8 shows that in the case of the old and ill grasshopper, all five subjects disagreed more or less strongly with

<sup>18</sup>This divergence illustrates the danger of ethical factors which can be read in more than one way.

the ant's decision, as shown by the negative values for question 3. However, in the case of the mocking scenario, four of the five subjects tested agreed with the ant's decision. This would suggest that previous information is being used to evaluate not just the players in this drama, but also the decisions being made.

Note that Subject 5 runs counter to the positive evaluation of the ant's decision. Remember that this was the subject who based the evaluation of the grasshopper on his singing. It may be that the negative evaluation of the ant's decision represents a reaction to the positive model of the grasshopper that this subject has just indicated.

### 7.3 Ethical decisions and their consequences

As we noted above, the test environment used here spells out explicitly the fatal consequences for the grasshopper of the ant's ethical judgment. In the case of Figure 8, we can see a sharp divergence between three subjects who saw the mocking grasshopper's death as quite inappropriate (with scores of -6 to -10) and two subjects who saw the death as at least somewhat appropriate (with scores of +2 and +3).

We find a similar divergence in Figure 8, where four subjects see the old and ill grasshopper's death as inappropriate, while one (Subject 1), sees the death as at least marginally appropriate, with a score of +1. Note that this subject has also given the grasshopper the lowest initial score and that the curve of this subject's judgments never diminishes over the course of the scenario. This divergent perspective merits further attention. In general, though, the data presented appears to show that the gravity of the consequences of some act has a measurable effect on reader judgments. This could be tested by having, for example, the ant cause the grasshopper to miss a single meal.

### 7.4 From ethical paragon to ethical agent

As noted above, the ant is at least partially responsible for the grasshopper's death from starvation. It is thus possible to ask how the ant itself is judged. Consideration of the two Figures shows that in all cases but one (Subject 4, whose evaluation of the grasshopper's demise was positive but of the ant negative) subjects who considered the grasshopper's demise to be inappropriate judged the ant's behaviour negatively, while those who saw the grasshopper's death to be appropriate judged the ant's behaviour to be positive. Note that this correlation traverses the two situations. Thus, even in the *old and ill* scenario, Subject 1 is consistent in answering Questions 4 and 5, as is Subject 5 in the *mocking* scenario, by evaluating both the demise and the ant's decision as positive.

This consistency would suggest that readers are capable of focusing on particular players in a scenario as they become involved with ethical decisions.

### 7.5 Correlations

The small number of cases considered here, as well as the relatively primitive scenarios used, preclude any strong conclusions. However, the data obtained does provide at least initial confirmation that narrative order can inflect ethical judgments, that the evaluation of characters appears to contribute to the evaluation of outcomes, and that the gravity of outcomes may play a role in ethical considerations. Clearly, further work would be required to test these hypotheses in a more robust fashion.

## 8 Applications of an Ethics Workbench

As we have seen, the interactions between narrative and ethical evaluations are multi-faceted, complex and, to date, little studied or understood. We believe that, properly extended, the proof-of-concept framework presented here has the potential to contribute to research in several areas:

- (i) it provides a locus of discussion for dealing with ethical issues, for example in Philosophy classes. Instead of discussing principles in the abstract, it would provide a place where principles could be instantiated in stories, either by examining the outcomes generated by different weightings of ethical principles, or by studying reader reactions to developing scenarios. By varying these along one or more dimensions, it would be possible to elicit discussion of the role of the various factors. For example, consider an undergraduate class in which the Ant and Grasshopper story is retold in terms of human beings. Would subjects still be willing to see a *lazy person* die, as opposed to an insect?<sup>19</sup> Or what about the situation where the ant has a lot of food versus the one where he or she has little, and perhaps a large family?
- (ii) it provides a testbed for studying, in considerably more detail than we have done here, the interplay of narrative and reader judgments. Along with changes of ordering, it would be possible to study the effect of particular elements on judgments. For example, what are the elements which set a positive or negative polarity?
- (iii) it would provide a means of ‘crowdsourcing’ reader judgments, by the capture of evaluations across many readers. At a finer level, it might permit the identification of cultural or other factors which affect narratives. For example, do men and women tend to react the same way to the ant and the grasshopper story? Do members of different political stripes or religions react differently? At the extreme, an environment like this provides a means of ‘shopping’ political or even marketing narratives among target groups.
- (iv) finally, to the extent that reader judgments are allowed to inflect narratives, an environment like this provides one element of a dynamic narrative structure, a sort of ‘choose your own narrative’ with an ethical dimension.

## 9 Conclusions and Next Steps

The work presented here is still very much a preliminary exploration of a new area in the computational modelling of narrative. It has necessarily involved the use of a number of simplifying assumptions. Among other things, we have assumed that humans are capable of expressing their moral judgments by means of a scalar representation, that these scales are in fact linear, and that the combination of more than one scale is additive. Despite these caveats, we believe that the use of DAGs, combined with an interface which presents text as ordered discrete elements and a means of obtaining reader judgments, has the potential to open windows into the interplay between reader and text.

In previous work, we have assumed that, since texts are massively underspecified (REF Pavel), the interplay of reader and text is particularly important. We have also assumed that the contents of a reader’s mind may be represented by a DAG and that reading represents a complex situation of DAG-synthesis. Although we have shown ethical principles in this

---

<sup>19</sup>The parallels between the Aesop text and Dicken’s *Christmas Carol* are striking. “Are there no workhouses? Are there no prisons?”

paper as individual items, or lists, it is far more likely that they have complex interactions among each other and that they also may be represented by DAGs. This will require further exploration.

---

#### References

---

- 1 Wayne C. Booth. *The company we keep: an ethics of fiction*. University of California Press, Berkeley, 1988.
- 2 Brian Boyd. *On the origin of stories: evolution, cognition, and fiction*. Belknap Press of Harvard University Press, Cambridge, Mass., 2009.
- 3 Carol Gilligan. *In a different voice: psychological theory and women's development*. Harvard University Press, Cambridge, Mass., 1982.
- 4 Michael Levison, Greg Lessard, Craig Thomas, and Matthew Donald. *The Semantic Representation of Natural Language*. Bloomsbury, London, 2012.
- 5 Thomas G. Pavel. *Fictional worlds*. Harvard University Press, Cambridge, Mass., 1986.
- 6 James Phelan. Narrative judgments and the rhetorical theory of narrative: Ian McEwan's *Atonement*. In James Phelan and Peter J. Rabinowitz, editors, *A companion to narrative theory*, pages 322–336. Blackwell, Oxford, 2005.
- 7 Jane Stadler. *Pulling focus: intersubjective experience, narrative film, and ethics*. Continuum, New York, 2008.