

Role of Kolmogorov Complexity on Interest in Moral Dilemma Stories

Antoine Sillienfest, Jean-Louis Dessalles

({antoine.sillienfest, jean-louis.dessalles}@telecom-paristech.fr)
Telecom ParisTech, INFRES, 46 Rue Barrault, 75013, Paris

Abstract

Several studies have highlighted the combined role of emotions and reasoning in the determination of judgments about morality. Here we explore the influence of Kolmogorov complexity in the determination, not only of moral judgment, but also of the associated narrative interest. We designed an experiment to test the predictions of our complexity-based model when applied to moral dilemmas. It confirms that judgments about interest and morality may be explained in part by discrepancies in complexity. This preliminary study suggests that cognitive computations are involved in decision-making about emotional outcomes.

Keywords: Kolmogorov complexity; moral dilemma; moral judgment; narrative interest; emotion.

Introduction

Humans devote a considerable amount of time to producing narratives. Spontaneous conversational narratives constitute a large amount of our conversational time (Norrick, 2000) and fictional narratives constitute a large part of human productions (*e.g.* novels, movies, video games). Modeling narrative interest and emotional impact in narratives, especially in fictional narratives, is of major importance, both scientifically and economically, as there are a variety of potential applications (*e.g.* serious games, film industry, video games). The selection of events that people consider relevant to tell is a not yet fully understood process. Only a small proportion of our experiences passes the selection. Moral dilemma belong to the situations that make good stories.

Previous studies have pointed out that emotional intensity and complexity drop have a decisive influence on narrative interest. The aim of this article is to explore the role of complexity change in the determination of morality and interest in moral dilemma.

Morality judgments in dilemma

The number of studies on morality and emotion grew steadily in the 1980s and 1990s, and even more during the last decade. Various disciplines now investigate human morality and the interplay of emotion and reason in moral judgment and decision-making.

Moral psychology initially focused on reasoning. During the 1950s and 1960s, mental models and information processing were the preferred framework in psychology. Kohlberg (1958) proposed a six-stage developmental model of moral reasoning which, he thought, drives moral judgment. In the 1980s, however, the idea that moral emotions also play a role has been highlighted. New findings in evolutionary psychology and in primatology pointed to the crucial role of a specific set of emotions. This "affective revolution" has been reinforced during the two last decades. Recent evidence suggests

that moral judgment is more a matter of emotion and affective intuition than of deliberate reasoning (Haidt & Hersh, 2001). Emotion now plays a central role in moral psychology research (Haidt, 2007). Evidence that emotions guide moral judgments comes from brain imagery (Moll et al., 2002; Greene et al., 2004; Moll & Oliveira-Souza, 2007; Koenigs et al., 2007; Decety, Michalska & Kinzler, 2011), philosophy (Roeser, 2006), and psychology (Wheatley & Haidt, 2005; Valdesolo & DeSteno, 2006; Schnall et al., 2008).

Recently, new findings from several areas of cognitive neuroscience have suggested that emotions and reasoning both matter, but that automatic emotional processes tend to dominate (Greene & Haidt, 2002; Greene et al., 2004).

Various elements of particular importance to our study, such as social consensus, proximity (the feeling of nearness), the magnitude of consequences or the probability of effect have been shown to affect our judgment in moral dilemma (Jones, 1991). It has also been shown that people judge permissible to harm people as a side effect but not as a means (Cushman, Young & Hauser, 2006).

Interest in narratives

Both reasoning and emotions seem to control the intensity of narrative interest. It has been observed that human cognition is sensitive to complexity, in the sense of Kolmogorov (*i.e.* the length of the minimal determination of a situation) (Chater, 1999; Chater & Vitényi, 2003). Simplicity Theory (Dessalles, 2008a; see also www.simplicitytheory.org) highlighted the role of unexpectedness in the selection of interesting events: a situation is unexpected if it is more complex to produce than to determine. This means that the generation complexity C_w of an unexpected event is higher than the complexity C_d of its determination.

C_w measures the minimum quantity of information that must be given for the "world" (as the observer knows it) to make the situation happen. It evaluates the size of the minimal explication of the situation. C_d measures the quantity of information needed by the observer to describe the situation unambiguously. It evaluates the size of the minimum description of the situation.

Unexpectedness U is the difference between the generation complexity and the determination complexity ($U = C_w - C_d$). The study of unexpectedness makes good predictions about which parameters control narrative interest in situations such as fortuitous encounters, atypical events, coincidences or rare events (Dessalles, 2008b; see also www.simplicitytheory.org).

Interest is also a matter of emotions. Emotional intensity plays a crucial role in the selection of narrated events (Rimé,

2005). In what follows, we will present a model in which emotional intensity results from the combination of the emotional category of the event (whether it is about ten deaths or merely about a ten-Euro loss) with unexpectedness. Then we will present an experiment designed to test the model. Finally, we will discuss the validity and the generality of this approach based on complexity.

A complexity-based model of narrative interest and of moral judgment

The model proposed in this article intends to show that variations of Kolmogorov complexity may contribute to explain both moral judgment and interest in moral dilemma stories.

Any outcome i comes with an hypothetical emotional intensity $E_h^i (> 0)$ attached to it. This value does not take the valence (positive or negative) of the emotion into account, but only the standard magnitude of the corresponding class of events (sometimes considered to result from social consensus) (e.g. the death of child is supposed to be emotionally more intense than the death of an adult person, all things being equal) (Jones, 1991; Bleske-Rechek et al. 2010).

According to Simplicity Theory, the emotional intensity attached to an event is the sum of the hypothetical emotional intensity and of unexpectedness: $E = E_h + U$. Since U depends on the complexity of the persons involved in the event (with a minus sign), the definition of E reflects the fact that one is more affected if the victim of an accident is a close acquaintance or a celebrity (as close acquaintances or celebrities require less information to be determined).

When considering the narrative value of an event s which is the outcome of an action a , the computation of E is performed *ex post*: E is derived from E_h and U . In *ex post* calculus, $C_w(s) = C_w(s|a) + C_w(a)$ in which $C_w(s|a)$ measures the amount of information the world needs to produce s from a (see Figure 1). To judge the moral value of a , the computation is *ex ante*: s is evaluated from its emotional intensity E , from the causal unexpectedness $U(s|a) = C_w(s|a) - C_d(s|a)$ and from the unexpectedness of action a . $C_d(s|a)$ measures what is still to be determined about s once a is known. Then the moral evaluation E_h^a of a is $E_h^a(s) = E(s) - U(s|a) - U(a)$.

We introduce the notion of responsibility: $R^a(s) = C_w(s) - C_w(s|a)$. The more a makes s easy (resp. hard) to produce, the more $R^a(s)$ increases (resp. decreases) and the more (resp. less) the actor is judged responsible for s . We also introduce the notion of targetting: $T^a(s) = C_d(s) - C_d(s|a)$ which evaluates the contribution of a in the description of s . $T_a(s) = C_d(s)$ means that s is fully described by a , the outcome s is targeted by the actor that does a . Lastly, we introduce the notion of inadvertence $F^a = U(a)$ which measures how unexpected the action is. If F^a is large, then a has been done inadvertently. Eventually, $E_h^a = R^a - T^a - F^a$.

In this paper, we only consider premeditated actions ($F^a = 0$) that do fully describe the outcomes ($C_d(s|a) = 0$). Therefore:

$$E_h^a(s) = E(s) - C_w(s|a)$$

We define the emotional gain ΔE_a of a moral dilemma as the difference between emotional intensities for the desired consequences and undesired consequences (see Figure 1). ΔE_a estimates how satisfying the consequences of an action appear (note that ΔE_a can be negative or positive). The model leads us to the following predictions:

1. The narrative interest I of a situation increases with its unexpectedness U of its hypothetical emotional intensity E_h .

$$I = E_h + U$$

2. The moral judgment MJ , in the case of an action, increases with the emotional gain ΔE_a .

$$MJ = \Delta E_a$$




Figure 1: Complexity and Emotional intensity in a cause-to-consequence schema

The model leads to the following particular expectations:

1. A character's actions that have negative outcomes will be less morally approved if they are more direct, because their causal complexity is smaller.
2. The complexification of a causal chain will increase the narrative interest of an event, because it increases the generation complexity of the consequences, which thus appear more unexpected.
3. An action will appear more interesting, but will be less approved, if its negative consequences are simpler. In particular, an action that provokes the death of relatives or family members raises more interest (after the fact) but will be less approved.
4. Unexpected events that alter the normal course of a causal chain will have a positive influence on narrative interest.

Experiment

Participants were asked to take the perspective of a reader and to evaluate how alternative endings of a moral dilemma story would be globally perceived by other readers on two aspects: narrative interest and moral approval of the character's action. Participants therefore were not supposed to engage their own judgment.

Participants had to read the following story (original in French):




Figure 2: Illustration for the *flooded mine* dilemma

Tom is a miner who left Scotland, his native country, to go to work in a mine in Argentina. Here is what happened to him two years ago, while he was watching the mining site upstream. The river flooded suddenly. It was flooding one of the two tunnels of the mine. Tom knew that there were five people in this tunnel (tunnel A). In the other tunnel (tunnel B), there was only one person. The water level was rising fast, and the current in the tunnel A was increasing. The trapped persons were going to drown. Tom knew that by interrupting the current in one tunnel, one would flood the other one and provoke the death of people in it. The current cannot be interrupted in both tunnels at the same time. Tom was the only one who could act. He stood at the entrance of the two tunnels, near a crane and a heavy and voluminous box.

This *flooded mine* dilemma (Figure 2) is largely inspired from the classical trolley problem (Thomson, 1985). The context has been augmented to allow a larger variety of alternative endings.

Method, Participants and Materials

A total of 64 individuals (aged from 19 to 65 y.o., mean 26.11 (std. dev. 7.72), 26 females) participated to the test. The study was conducted online. Participants, mainly engineering students, were recruited via social networks and billposting.

They were asked to read the flooded mine story carefully. For each of the 4 phases of the test, several alternative endings were proposed, in which actions, causal links and consequences were varied. Alternatives were presented in the same order to all participants. Using the numbering defined later in this paper, the original orders for phases 1 to 4 were 1-2-3, 3-1-4-2, 2-4-1-3 and 1-4-2-3.

For each alternative, participants were asked to answer the two following questions on a 10-point scale:

1. "According to you, will the readers of the story approve Tom's actions?" (-5: "Disapprove", 5: "Approve")
2. "According to you, will the readers of the story find the alternative interesting?" (-5: "Not Interesting", 5: "Interesting")

We omitted the zero from this scale to force participants to choose between approval or disapproval (resp. interesting or not interesting). The answering times (with standard deviation) for the 4 phases were 3'32"(1'12"), 3'12"(1'36"), 1'36"(1'17") and 2'17"(1'31"). We manually checked the answer files for individuals who provided random or uncomplete results.

The different phases of the test successively explore the role that Kolmogorov complexity plays both on interest and on moral judgment.

Phase 1

Former studies have shown that harming actions are more likely to be judged moral if their consequences are more indirect (Cushman, Young, & Hauser, 2006). In phase 1 of the test, we tried to reproduce this result and explore how causal complexity also affects narrative interest. We also investigated the role of the complexity of the action.

Alternative endings of phase 1

1. Tom pushed the box in front of the entrance of the tunnel A in order to interrupt the current in this tunnel.
2. Tom got in the crane, grabbed the box with the crane's hook, brought the box above the entrance of the tunnel A and dropped it in the middle of the current to stop the current in the tunnel.
3. Tom broke the dam of tunnel B to flood this tunnel immediately.

It interrupted the current in the tunnel A, the tunnel B was flooded.

Five persons were saved, one person died by drowning

Results There is a main effect of intention (F-test: $F(1, 190) = 60.87, p < 0.0001$) on morality but no significant effect on interest ($F(1, 190) = 0.0045, p = 0.95$). A pairwise comparison of the two alternatives involving the box revealed no effect of the way the box is carried in the middle of the current on morality ($p = 1$) and interest ($p = 1$) (see Figure 3). In this phase, we could replicate a classical result of the trolley problem, in which people approve actions that lead to harming a victim as a side-effect but not as a means.

The main result of this first phase is that situations which are less approved by participants are not necessarily more interesting. Elements such as harming someone as a side effect or as a means only affect moral judgment. In this test, the emotional intensity of consequences is not manipulated. In alternatives 1 and 2, the causal chain between Tom's action and its harming consequence is more complex than in alternative 3; our model predicts that alternative 3 will be less approved than alternatives 1 and 2. Since all consequences are equally (un)expected, our model predicts no effect on interest.

More generally, our model predicts that actions are more likely to be approved if their positive effects are more direct and if their negative effects are more indirect.




Figure 3: Results of phase 1: mean approval (left panel) and interest (right panel) ratings for three alternatives in which the proposed harmful actions vary in their intentional status (more or less direct action on the victim). Error bars indicate the 95% confidence interval.

Phase 2

Phase 2 explores how the *length* of the causal chain of events between a causal action and its consequences affects moral judgments and narrative interest.

Alternative endings of phase 2

Tom pulled the box in front of tunnel A.

1. As he expected, it stayed across the current because of its weight.
2. It was carried by the current and, as Tom expected, it was stopped by the struts in the tunnel.
3. It was carried by the current, hit the struts in the tunnel and, as Tom expected, some struts got broken and part of the ceiling at the entrance of the tunnel collapsed.
4. It was carried by the current and hit the struts in the tunnel; beams fell down from the ceiling; they were also carried by the current and were stopped by other struts. As Tom expected, it formed a new dam.

It interrupted the current in tunnel A. The five persons were saved, but the tunnel B got flooded and one person died by drowning

Results There is a main effect of the length of the causal chain of events on both morality ($F(3, 252) = 3.01, p = 0.03$) and interest ($F(3, 252) = 3.29, p = 0.02$) (see Figure 4)

In this phase, only causal generation complexity is manipulated. Since unexpectedness U in I is an increasing function of this complexity, longer deterministic chain of events make the outcomes more unexpected. ΔE_a is a decreasing function of generation complexity for desired consequences. Our model correctly predicts that the outcomes will appear more interesting and that the actions will be less approved.

Jones (1991) used the expression *probability of effect* to refer to the probability that a harming event will occur. We suggest that the term of probability is not adapted, because in




Figure 4: Results of phase 2: mean approval (left panel) and interest (right panel) ratings for four alternatives in which the proposed causal chain of events vary in their length ($L1 < L2 < L3 < L4$). Error bars indicate the 95% confidence interval.

many cases, the length of the chain of events is what is relevant, even if it is deterministic. The notion of unexpectedness correctly captures this phenomenon.

Phase 3

Previous studies on morality (Bleske-Rechek et al., 2010) have shown that the identity of victims may affect our moral judgment. This phase explores how the identity of the single victim in tunnel B influences both moral judgment and narrative interest.

Alternative endings of phase 3

Tom knew that [#1] was in tunnel B. Tom pushed the box in front of the entrance of tunnel A. It interrupted the current in tunnel A, tunnel B was flooded. Five persons were saved, [#2] died.

[#1]/[#2] were :

1. someone / one person
2. one of his friends / Tom's friend
3. his own cousin / Tom's cousin
4. a 10-years old child / the child

Results There is a main effect of the identity of the victim (undefined person, cousin, friend and 10-year old child) on morality ($F(3, 252) = 11.79, p < 0.05$) and interest ($F(3, 252) = 14.26, p < 0.05$) ratings. A series of pairwise contrasts clarifies the nature of this interaction. The presence of an undefined person elicited significantly higher moral approval ratings (*undefined person* vs. *friend*, *cousin*, *child*, respectively: $F(1, 126) = 11.32, 14.07$ and $38.25; p = 0.001, 0.0003$ and < 0.0001) and significantly lower interest ratings (*undefined person* vs. *friend*, *cousin*, *child*, respectively: $F(1, 126) = 34.37, 18.32$ and $24.87; p < 0.0001, < 0.0001$ and < 0.0001) (see Figure 5).

Alternatives in which victims are less complex to describe (Tom's cousin and Tom's friend) are more interesting than the ones involving some undefined victim. As suggested by our




Figure 5: Results of phase 3: mean approval (left panel) and interest (right panel) ratings for four alternatives in which the victims vary in their identity. Error bars indicate the 95% confidence interval.

model, the unexpectedness increases and the emotional gain decreases when undesired consequences are simpler. The model correctly predicts that alternatives 2 and 3 will be less approved and appear more interesting than alternative 1.

The fact that the victim is a child increases the hypothetical emotional intensity. Our model correctly predicts that harming a child will be less approved but more interesting.

Our model agrees with results of Bleske-Rechek et al. (2010) and with some aspects of Jones' "components of moral intensity" such as proximity (defined as the feeling of nearness the moral agent has for the victim) and the social consensus (defined as a degree of social agreement that a proposed act is evil) (Jones, 1991).

Phase 4

Two elements are manipulated in phase 4: the more or less direct action of Tom and the presence or absence of unexpectedness in the course of events.

Alternatives

- Tom broke the dam of tunnel B to flood this tunnel immediately. What happened was not expected. The fragments of the dam formed a new dam at the entrance of tunnel B. It was not enough to stop the current in tunnel A. Five persons died, one person was saved.
- Tom broke the dam of tunnel B to flood this tunnel immediately. It interrupted the current in tunnel A, the tunnel B was flooded. Five persons were saved, one person died.
- Tom pushed the box in front of the entrance of the tunnel A in order to interrupt the current in this tunnel. What happened was not expected: the box was not big enough to stop the current in tunnel A. Five persons died, one person was saved.
- Tom pushed the box in front of the entrance of tunnel A in order to interrupt the current in this tunnel. It interrupted the current in tunnel A. Tunnel B was flooded. Five persons were saved, one person died.

Results There is a main effect of unexpectedness in the course of events for both morality ratings and interest ratings

(*non unexpectedness* vs. *unexpectedness*, respectively : $F(1, 252) = 14.66$ and 9.30 ; $p = 0.0002$ and 0.003). As in the phase 1, the action of Tom has significant effect on approval ratings (*push the box* vs. *break the dam*, $F(1, 252) = 16.03$, $p < 0.0001$) but no significant effect on interest ratings ($F(1, 252) = 1.83$, $p = 0.18$) (see Figure 6).




Figure 6: Results of phase 4: mean approval (left panel) and interest (right panel) ratings for four alternatives in which the proposed harmful actions vary in their intentional status and the courses of events vary in the presence or non presence of unexpectedness. Error bars indicate the 95% confidence interval.

These results are consistent with the observations of the previous phases. They show that variations of Tom's action only affect moral judgment. They also show that a combined increase of unexpectedness and of hypothetical emotional intensity positively influences interest but negatively affects our moral judgment about Tom's action.

This test does not only confirm previous results. It also confirms the prediction that the combination *breaking the dam + unexpected outcome* will be the least approved alternative (and the most interesting one), whereas the combination *pushing the box + deterministic result* will be the most approved alternative (and the least interesting one). However, the model is not accurate enough to predict which one among the alternatives 2 and 3 will be more approved or more interesting. Such ranking would require precise estimates of complexity values, which depend on observers' knowledge and personal history.

Discussion

The aim of this study was to present and test a complexity-based model for narrative interest and moral judgment in moral dilemma stories. The point was to explore the role played by Kolmogorov complexity. The model, based on Simplicity Theory, intends to offer a purely cognitive account of some judgements about morality and interest in moral dilemma stories, without using *ad hoc* hypothesis about moral or immoral classes of actions. This preliminary and exploratory study has shown that the model makes correct predictions and may explain classical results about moral judgements in terms of complexity. It intends to point out the dif-

ferent but complementary roles of complexity and emotional intensity in moral judgment and narrative interest.

Our moral judgments depend on the estimated emotional gain of the dilemma, which depends not only on the emotion attached to outcomes, but also on the complexity of the causal links that lead to them. A positive gain, for example an action which saves five but kills one, may be less morally approved if it appears that this positive consequences are more complex to produce than negative ones. The gain would also be considered as uncertain or unexpected. This contributes to explain why humans are highly sensitive to actions that would jeopardize the lives of relatives, friends, or family members.

Emotional intensity also plays an important role. This may be related to Haidt's work (Haidt & Hersh, 2001) that point out the role of automatic emotional processes. Due to cultural or societal consensus, some situations appears more emotionally intense than others.

Our model makes good qualitative predictions. In future work, we will explore how quantitative parameters that influence complexity (e.g. distance in space or time) affect the emotional intensity attached to outcomes. The identification and control of these quantitative parameters open the way to a variety of potential applications, for example in the domain of decision-making. Several factors that are spontaneously attributed to the emotional component of decision-making may be reinterpreted as complexity-based computations. Such an account, if valid, would not only be relevant from a scientific perspective. It would also be potentially useful to help decision-makers evaluate or anticipate certain decisions in which emotions are supposed to play a major role.

References

- Bleske-Rechek, A., Nelson, L., Baker, J., Remiker, M., & Brandt, S. (2010). Evolution and the trolley problem : people save five over one unless the one is young, genetically related, or a romantic partner. *Journal of Social, Evolutionary, and Cultural Psychology*, 4(3), 115–127.
- Chater, N. (1999). The search for simplicity: A fundamental cognitive principle ? *The Quarterly Journal of Experimental Psychology*, 52A(2), 273–302.
- Chater, N., & Vitányi, P. (2003). Simplicity : a unifying principle in cognitive science ? *TRENDS in Cognitive Sciences*, 7(1), 19–22.
- Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment. *Psychological Science*, 17(12), 1082–1089.
- Decety, J., Michalska, K. J., & Kinzler, K. D. (2011). The developmental neuroscience of moral sensitivity. *Emotion Review*, 3(3), 305–307.
- Dessalles, J.-L. (2008a). Coincidences and the encounter problem: A formal account. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th annual conference of the cognitive science society (cogsci 2008)* (pp. 2134–2139).
- Dessalles, J.-L. (2008b). *La pertinence et ses origines cognitives : nouvelles théories*. Hermes-Science Publications.
- Greene, J. D., & Haidt, J. (2002). How (and where) does moral judgment work ? *Trends in Cognitive Sciences*, 6(12), 517–523.
- Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron*, 44(2), 389–400.
- Haidt, J. (2007). The new synthesis in moral psychology. *Science*, 316, 998–1002.
- Haidt, J., & Hersh, M. A. (2001). Sexual morality: The cultures and emotions of conservatives and liberals. *Journal of Applied Social Psychology*, 31(1), 191–221.
- Jones, T. M. (1991). Ethical decision making by individuals in organizations: An issue-contingent model. *The Academy of Management Review*, 16(2), 366–395.
- Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., et al. (2007). Damage to the prefrontal cortex increases utilitarian moral judgements. *Nature*, 446, 908–911.
- Kohlberg, L. (1958). *The development of modes of moral thinking and choice in the years 10 to 16*. Unpublished doctoral dissertation, University of Chicago. (Unpublished)
- Moll, J., & Oliveira-Souza, R. de. (2007). Moral judgments, emotions and the utilitarian brain. *Trends in Cognitive Sciences*, 11(8), 319–321.
- Moll, J., Oliveira-Souza, R. de, Eslinger, P. J., Bramati, I. E., Mourão-Miranda, J., Andreiuolo, P. A., et al. (2002). The neural correlates of moral sensitivity: A functional magnetic resonance imaging investigation of basic and moral emotions. *The Journal of Neuroscience*, 22(7), 2730–2736.
- Norrick, N. R. (2000). *Conversational narrative: Storytelling in everyday talk*. John Benjamins Publishing Company.
- Rimé, B. (2005). *Le partage social des émotions*. Presses Universitaires de France.
- Roeser, S. (2006). The role of emotions in judging the moral acceptability of risks. *Safety Science*, 44(8), 689–700.
- Schnall, S., Haidt, J., Clore, G. L., & Jordan, A. H. (2008). Disgust as embodied moral judgment. *Personality and Social Psychology Bulletin*, 34(8), 1096–1109.
- Thomson, J. J. (1985). The trolley problem. *The Yale Law Journal*, 94(6), 1395–1415.
- Valdesolo, P., & DeSteno, D. (2006). Manipulations of emotional context shape moral judgment. *Psychological Science*, 17(6), 476–477.
- Wheatley, T., & Haidt, J. (2005). Hypnotic disgust makes moral judgments more severe. *Psychological Science*, 16(10), 780–784.