

Natural Language Generation enhances human decision-making with uncertain information

Dimitra Gkatzia
School of Computing
Edinburgh Napier University
Edinburgh, EH10 5DT, UK
d.gkatzia@napier.ac.uk

Oliver Lemon
Interaction Lab
Heriot-Watt University
Edinburgh, EH14 4AS, UK
o.lemon@hw.ac.uk

Verena Rieser
Interaction Lab
Heriot-Watt University
Edinburgh, EH14 4AS, UK
v.t.rieser@hw.ac.uk

Abstract

Decision-making is often dependent on uncertain data, e.g. data associated with confidence scores or probabilities. We present a comparison of different information presentations for uncertain data and, for the first time, measure their effects on human decision-making. We show that the use of Natural Language Generation (NLG) improves decision-making under uncertainty, compared to state-of-the-art graphical-based representation methods. In a task-based study with 442 adults, we found that presentations using NLG lead to 24% better decision-making on average than the graphical presentations, and to 44% better decision-making when NLG is combined with graphics. We also show that women achieve significantly better results when presented with NLG output (an 87% increase on average compared to graphical presentations).

1 Introduction

Natural Language Generation (NLG) technology can achieve comparable results to commonly used data visualisation techniques for supporting accurate human decision-making (Gatt et al., 2009). In this paper, we investigate whether NLG technology can also be used to support decision-making when the underlying data is uncertain. Current data-to-text systems assume that the underlying data is precise and correct – an assumption which is heavily criticised by other disciplines concerned with decision support, such as medicine (Gigerenzer and Muir Gray, 2011), environmental modelling (Beven, 2009), climate change (Manning et al., 2004), or weather forecasting (Kootval, 2008). However, simply presenting numerical ex-

pressions of risk and uncertainty is not enough. Psychological studies on decision making have found that a high percentage of people do not understand and can't act upon numerical uncertainty (Cokely et al., 2012; Galesic and Garcia-Retamero, 2010). For example, only 28% of Germans and 25% of Americans are able to answer the question: “Which of the following numbers represents the biggest risk of getting a disease: 1 in 100, 1 in 1000, 1 in 10?” (Galesic and Garcia-Retamero, 2010).

So far, the NLG community has investigated the conversion of numbers into language (Power and Williams, 2012) and the use of vague expressions (van Deemter, 2009). In this work, we explore how to convert numerical representations of uncertainty into Natural Language so as to maximise confidence and correct outcomes of human decision-making. We consider the exemplar task of weather forecast generation. We initially present two NLG strategies which present the uncertainty in the input data. The two strategies are based on (1) the World Meteorological Organisation (WMO) (Kootval, 2008) guidelines and (2) commercial forecast presentations (e.g. from BBC presenters). We then evaluate the strategies against a state-of-the-art graphical system (Stephens et al., 2011), which presents the uncertain data in a graphical way. Figure 1 shows an example of this baseline graphical presentation. We use a game-based setup (Gkatzia et al., 2015) to perform task-based evaluation, to investigate the effect that the different information presentation strategies have on human decision-making.

Weather forecast generation is a common topic within the NLG community, e.g. (Konstas and Lapata, 2012; Angeli et al., 2010; Belz and Kow, 2010; Sripada et al., 2005). Previous approaches have not focused on how to communicate uncertain information or the best ways of referring to




Figure 1: Graphics for temperature data.

probabilities of meteorological phenomena to occur. In addition, their evaluation is based on user ratings of grammatically, semantic correctness, fluency, coherence or via post-edit evaluation. Although these metrics are indicative of the quality of the text produced, they do not measure the impact the texts might have in people’s comprehension of uncertainty or on their ability to make decisions based on the information conveyed.

Our contributions to the field are as follows: (1) We study a principled mapping of uncertainty to Natural Language and provide recommendations and data for future NLG systems; (2) We introduce a game-based data collection environment which extends task-based evaluation by measuring the impact of NLG on decision-making (measuring user confidence and game/task success); and (3) We show that effects of the different representations vary for different user groups, so that user adaptation is necessary when generating multimodal presentations of uncertain information.

2 The Extended Weather Game

In this section, we present our extended version of the MetOffice’s Weather Game (Stephens et al., 2011). The player has to choose where to send an ice-cream vendor in order to maximise sales, given weather forecasts for four weeks and two locations. These forecasts describe (1) predicted rainfall (Figure 2) and (2) temperature levels together

Likelihood of occurrence	Lexicalisation
$p > 0.99$	“extremely likely”
$0.90 \leq p \leq 0.99$	“very likely”
$0.70 \leq p \leq 0.89$	“likely”
$0.55 \leq p \leq 0.69$	“probable - more likely than not”
$0.45 \leq p \leq 0.54$	“equally likely as not”
$0.30 \leq p \leq 0.44$	“possible - less likely than not”
$0.10 \leq p \leq 0.29$	“unlikely”
$0.01 \leq p \leq 0.09$	“very unlikely”
$p < 0.01$	“extremely unlikely”

Table 1: WMO-based mapping of likelihoods.

with their likelihoods in three ways: (a) through graphical representations (which is the version of the original game), (b) through textual forecasts, and (c) through combined graphical and textual forecasts. We generated the textual format using two rule-based NLG approaches as described in the next section. Users are asked to initially choose the best destination for the ice-cream vendor and then they are asked to state how confident they are with their choice. Based on their decisions and their confidence levels, the participants are finally presented with their “monetary gain”. For example, the higher the likelihood of sunshine, the higher the monetary gain if the player has declared that s/he is confident that it is not going to rain and it doesn’t actually rain. In the opposite scenario, the player would lose money. The decision on whether rain occurred is estimated by sampling the probability distribution. At the end of the game, users were scored according to their “risk literacy” following the Berlin Numeracy Test (Cokely et al., 2012). Further details are presented in (Gkatzia et al., 2015).

3 Natural Language Generation from Uncertain Information

We developed two NLG systems, WMO-based and NATURAL, using SimpleNLG (Gatt and Reiter, 2009), which both generate textual descriptions of rainfall and temperature data addressing the uncertain nature of forecasts.

WMO-based: This is a rule-based system which uses the guidelines recommended by the WMO (Kootval, 2008) for reporting uncertainty, as shown in Table 1. Consider for instance a forecast of sunny intervals with 30% probability of rain. This WMO-based system will generate the following forecast: “Sunny intervals with rain being possible - less likely than not”.

NATURAL: This system imitates forecasters and

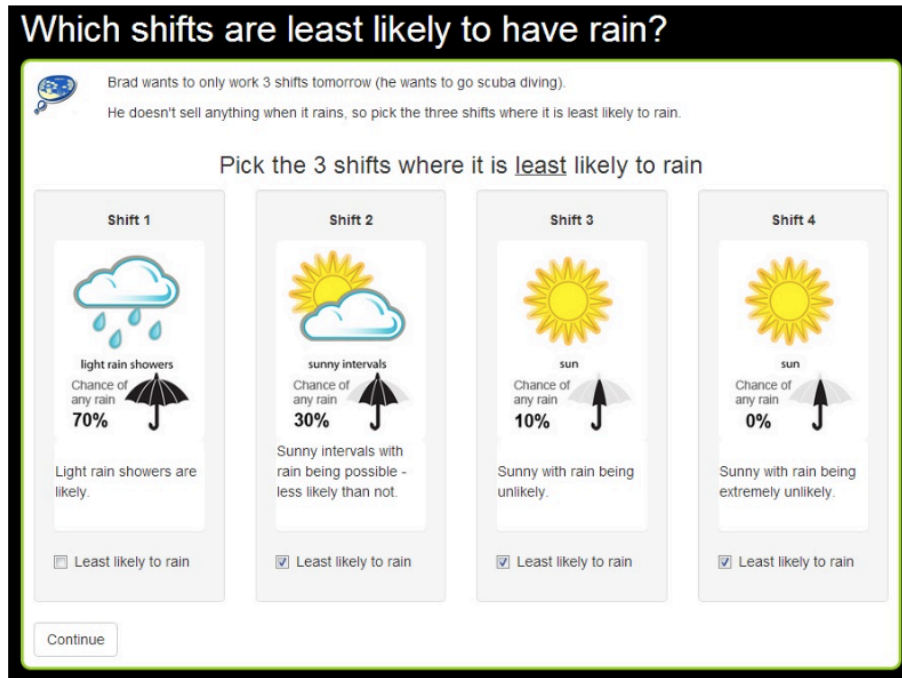


Figure 2: Screenshot of the Extended Weather Game (Rainfall: Graphics and WMO condition).

their natural way of reporting weather. The rules used in this system have been derived by observing the way that experts (e.g. BBC weather reporters) produce forecasts. For the previous example (sunny intervals with 30% probability of rain), this system will generate the following forecast: “Mainly dry with sunny spells”.

4 Evaluation

In order to investigate what helps people to better understand and act upon uncertainty in information presentations, we use five conditions within the context of the Extended Weather Game:

1. **Graphics only:** This representation shows the users only the graphical representation of the weather forecasts. For this condition we used the graphs that scored best in terms of human comprehension from (Stephens et al., 2011).
2. **Multi-modal Representations:**
 - **Graphics and NATURAL:** This is a multi-modal representation consisting of graphics (as described in the previous condition) and text produced by the NATURAL system.
 - **Graphics and WMO-based:** This is also a multi-modal representation consisting of graphics and text produced by the WMO-based system.

3. NLG only:

- **NATURAL only:** This is a text-only representation as described above.
- **WMO-based system only:** This is also a text-only representation.

5 Data

We recruited 442 unique players (197 females¹, 241 males, 4 non-disclosed) using social media. We collected 450 unique game instances (just a few people played the game twice). The anonymised data will be released as part of this submission.

6 Results

In order to investigate which representations assist people in decision-making under uncertainty, we analysed both the players’ scores (in terms of monetary gain) and their predictions for rainfall with regard to their confidence scores. As we described in Section 2, the game calculates a monetary gain based on both the decisions and the confidence of the player, i.e. the decision-making ability of the player. Regarding confidence, we asked users to declare how confident they are on a 10-point scale. In our analysis we therefore focus on both confidence and score at the game.

¹Women made up 44.5% of the subjects.

	Monetary gains	Confidence
Graphs only	81.15	78.5%
Multi-modal	117.51	83.7%
NLG only	101.33	66%

Table 2: Average Monetary gains and Confidence scores (All Adults).

6.1 Results for all adults

Multi-modal vs. Graphics-only: We found that use of multi-modal representations leads to gaining significantly higher game scores (i.e. better decision-making) than the Graphics-only representation ($p = 0.03$, effect = +36.36). This is a 44% average increase in game score.

Multi-modal vs. NLG-only: However, there is no significant difference between the NLG only and the multi-modal representation, for game score.

NLG vs. Graphics-only: We found that the NLG representations resulted in a 24.8% increase in average task score (i.e. better decision-making) compared to the Graphics-only condition, see Table 2: an average score increase of over 20 points. There was no significant difference found between the WMO and NATURAL NLG conditions.

Confidence: For confidence, the multi-modal representation is significantly more effective than NLG only ($p < 0.01$, effect = 17.7%). However, as Table 2 shows, although adults did not feel very confident when presented with NLG only, they were able to make better decisions compared to being presented with graphics only.

Demographic factors: We further found that prior experience on making decisions based on risk, familiarity with weather models, and correct literacy test results are predictors of the players’ understanding of uncertainty, which is translated in both confidence and game scores. In contrast, we found that the education level, the gender, or being native speaker of English does not contribute to players’ confidence and game scores.

6.2 Results for Females

We found that females score significantly higher at the decision task when exposed to either of the NLG output presentations, when compared to the graphics-only presentation ($p < 0.05$, effect = +53.03). This is an increase of 87%, also see Table 3. In addition, the same group of users scores significantly higher when presented with the multi-modal output as compared to graphics only ($p = 0.05$, effect = 60.74%). Interestingly, for

	Monetary gains	Confidence
Graphs only	60.83	74.6%
Multi-modal	118.41	81.3%
NLG only	113.86	65.8%

Table 3: Average Monetary gains and Confidence scores (Females).

this group, the multi-modal presentation adds little more in effectiveness of decision-making than the NLG-only condition, but the multi-modal presentations do enhance their confidence (+15%). We furthermore found that educated (i.e. holding a BSc or higher degree) females, who also correctly answered the risk literacy test, feel significantly more confident when presented with the multi-modal representations than with NLG only ($p = 0.01$, effect = 16.7%).

6.3 Results for Males

We found that males obtained similar game scores with all the types of representation. This suggests that the overall improved scores (for All Adults) presented above, are largely due to the beneficial effects of NLG for women. In terms of confidence, males are more likely to be more confident if they are presented with graphics only (81% of the time) or a multi-modal representation (85% of the time) ($p = 0.01$).

7 Conclusions and Future Work

We present results from a game-based study on how to generate descriptions of uncertain data – an issue which so far has been unexplored by data-to-text systems. We find that there are significant gender differences between multi-modal, NLG, and graphical versions of the task, where for women, use of NLG results in a 87% increase in task success over graphics. Multimodal presentations lead to a 44% increase for all adults, compared to graphics. People are also more confident of their judgements when using the multimodal representations. These are significant findings, as previous work has not distinguished between genders when comparing different representations of data, e.g. (Gatt et al., 2009). It also confirms research on gender effects in multi-modal systems, as for example reported in (Foster and Oberlander, 2006; Rieser and Lemon, 2008; Weiss et al., 2012). The results are also related to educational research, which shows that women perform better in verbal-logical tasks than visual-spatial tasks

(Zhu, 2007). An interesting investigation for future research is the interplay between uncertainty, risk-taking behaviour and gender, as for example reported in (Sarin and Wieland, 2016).

Acknowledgments

This research received funding from the EPSRC projects GUI (EP/L026775/1), DILiGENT (EP/M005429/1) and MaDrI-gAL (EP/N017536/1).

References

- Gabor Angeli, Percy Liang, and Dan Klein. 2010. A simple domain-independent probabilistic approach to generation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Anja Belz and Eric Kow. 2010. Extracting parallel fragments from comparable corpora for data-to-text generation. In *6th International Natural Language Generation Conference (INLG)*.
- Keith Beven. 2009. *Environmental Modelling: An Uncertain Future?* Routledge.
- Edward T. Cokely, Mirta Galesic, Eric Schulz, Saima Ghazal, and Rocio Garcia-Retamero. 2012. Measuring risk literacy: The berlin numeracy test. *Judgment and Decision Making*, 7(1):25–47.
- Mary Ellen Foster and Jon Oberlander. 2006. Data-driven generation of emphatic facial displays. In *Proc. of the 11th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Mirta Galesic and Rocio Garcia-Retamero. 2010. Statistical numeracy for health: A cross-cultural comparison with probabilistic national samples. *Archives of Internal Medicine*, 170(462–468).
- Albert Gatt and Ehud Reiter. 2009. SimpleNLG: A realisation engine for practical applications. In *ENLG*.
- Albert Gatt, Francois Portet, Ehud Reiter, James Hunter, Saad Mahamood, Wendy Moncur, and Somayajulu Sripada. 2009. From Data to Text in the Neonatal Intensive Care Unit: Using NLG Technology for Decision Support and Information Management. *AI Communications*, 22: 153–186.
- G. Gigerenzer and J. A. Muir Gray, editors. 2011. *Better doctors, better patients, better decisions: Envisioning health care 2020*. Cambridge MIT Press.
- Dimitra Gkatzia, Amanda Cercas Curry, Verena Rieser, and Oliver Lemon. 2015. A game-based setup for data collection and task-based evaluation of uncertain information presentation. In *Proceedings of the 15th European Workshop on Natural Language Generation (ENLG)*, pages 112–113, Brighton, UK, September. Association for Computational Linguistics.
- Ioannis Konstas and Mirella Lapata. 2012. Unsupervised concept-to-text generation with hypergraphs. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Haleh Kootval, editor. 2008. *Guidelines on Communicating Forecast Uncertainty*. World Meteorological Organisation.
- Martin Manning, Michel Petit, David Easterling, James Murphy, Anand Patwardhan, Hans-Holger Rogner, Rob Swart, and Gary Yohe. 2004. IPCC Workshop on Describing Scientific Uncertainties in Climate Change to Support Analysis of Risk and of Options.
- Richard Power and Sandra Williams. 2012. Generating numerical approximations. *Computational Linguistics*, 38(1):113–134, March.
- V. Rieser and O. Lemon. 2008. Learning effective multimodal dialogue strategies from wizard-of-oz data: Bootstrapping and evaluation. *Proceedings of ACL*, pages 638–646.
- Rakesh Sarin and Alice Wieland. 2016. Risk aversion for decisions under uncertainty: Are there gender differences? *Journal of Behavioral and Experimental Economics*, 60:1 – 8.
- Somayajulu G. Sripada, Ehud Reiter, and Lezan Hawizy. 2005. Evaluation of an NLG system using post-edit data. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- Liz Stephens, Ken Mylne, and David Spiegelhalter. 2011. Using an online game to evaluate effective methods of communicating ensemble model output to different audiences. In *American Geophysical Union, Fall Meeting*.
- Kees van Deemter. 2009. Utility and language generation: The case of vagueness. *Journal of Philosophical Logic*, 38(6):607–632.
- Benjamin Weiss, Sebastian Möller, and Matthias Schulz. 2012. Modality preferences of different user groups. In *The Fifth International Conference on Advances in Computer-Human Interactions (ACHI)*.
- Zheng Zhu. 2007. Gender differences in mathematical problem solving patterns: A review of literature. *International Education Journal*, 8(2):187 – 203.