

Taking decisions in a Hybrid Conversational AI architecture using Influence Diagrams

Roberto Basile Giannini^{1,2,†}, Antonio Origlia^{1,2,*,†} and Maria Di Maro^{1,2,†}

¹Dept. of Electrical Engineering and Information Technology, University of Naples Federico II

²Urban/ECO Research Center, University of Naples Federico II

Abstract

This paper explores the application of the Influence Diagrams model for decision-making in the context of conversational agents. The system consists of a *Conversational Recommender System* (CoRS), in which the decision-making module is separate from the language generation module. It provides the capability to evolve a *belief* based on user responses, which in turn influences the decisions made by the conversational agent. The proposed system is based on a pre-existing CoRS that relies on Bayesian Networks informing a separate decision process. The introduction of Influence Diagrams aims to integrate both Bayesian inference and the dialogue move selection phase into a single model, thereby generalising the decision-making process. To test the effectiveness and plausibility of the dialogues generated by the developed CoRS, a dialogue simulator was created and the simulated interactions were evaluated by a pool of human judges.

Keywords

Conversational AI, Decision-making, Influence Diagrams

1. Introduction

In recent years, the success of neural networks has generated significant enthusiasm among professionals in the field of artificial intelligence as well as the general public. Various applications, such as speech recognition, computer vision and even interactive conversational models like ChatGPT, have increasingly engaged users, inevitably shaping their perception of AI. This perception can have various implications, even within the scientific community. Attributing human-level intelligence to the tasks currently accomplished by neural networks is questionable, as these tasks barely rise to the level of abilities possessed by many animals [1]. Neural-based approaches to artificial intelligence have been criticised because of the limitations that are intrinsic to purely associative methods. One notable analysis of the problems that come when considering linguistic material generated without a real understanding of the *meaning* of what is being said is found in [2], which highlights that, because of the way it is generated, content produced by GPT models adheres to at least one formal definition of *bullshit*. The fundamental problem with these models is that, while they are trained to capture surface aspects of communication, they are never exposed to the *reasons why* language is produced. When they output the most

probable continuation of the provided prompt, they leave entirely to the human reader the task of interpreting what the produced output *might have meant*.

From a linguistics point of view, within the framework of Austin’s speech act theory [3] “saying something” equals “doing something”; the act of producing a sentence (*locutive act*) is fuelled by an intention (*illocutive act*) that produces changes in the world (*perlocutive act*). This classic view of the act of speaking highlights that conversation is a form of intervention in the world: it is put in action to alter in some way the conversational context. This same position is also found in the recent literature about the role of causality in artificial intelligence. Judea Pearl’s Ladder of Causation [4] puts *intervention* capabilities on the second level of the ladder, characterised by the verb *doing*, as in Austin’s seminal work. In this work, machine learning capabilities are limited to the first step of the Ladder, concerned with *observational* capabilities, leaving *interventional* ones out.

From this perspective, a conversational agent that produces language *motivated* by the achievement of a goal, thus modelling a *raison d’exprimer*, is an agent capable of using language with *interventional* purposes, which can be placed on the second step of the Ladder of Causation. A tool that aims to define conversational agents according to this philosophy is the Framework for Advanced Natural Tools and Applications with Social Interactive Agents (FANTASIA) [5], an Unreal Engine¹ plugin designed to develop embodied conversational agents. Built upon the functionalities offered by the tool, the FANTASIA Interaction Model follows these main principles: **Behaviour Trees** (BT) [6] are used to organise and pri-

CLiC-it 2024: Tenth Italian Conference on Computational Linguistics, Dec 04 – 06, 2024, Pisa, Italy

*Corresponding author.

[†]These authors contributed equally.

✉ ro.basilegiannini@studenti.unina.it (R. Basile Giannini);
antonio.origlia@unina.it (A. Origlia); maria.dimaro2@unina.it
(M. Di Maro)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



¹<https://www.unrealengine.com/>

oritise dialogue moves; **Graph Databases** (i.e., Neo4j [7]) are used for knowledge representation and dialogue state tracking; **Probabilistic Graphical Models** (PGM) are used for decision making; LLMs are used to verbalise the decisions taken by PGMs.

The latest results obtained using FANTASIA, presented in [8, 9], used a decision system based on Bayesian Networks to estimate probability distributions over ratings for users of a movie recommender system. The decision about the dialogue move was taken by a rule-based system taking into account these estimates. In this work, we further develop the approach by generalising the decision process using a single model, an Influence Diagram (ID) [10]. IDs represent an extension of BNs [11] since, in addition to probabilistic nodes, they also contain:

- *Decision nodes*, which represent decision points for the agent and which may be multiple within the model.
- *Utility nodes*, which represent utility (or cost) factors and which will drive the agent’s decisions, since the objective will be to maximise the utility of the model.

Consequently, in addition to the modelling of probabilistic inference problems, the use of IDs also enables the modelling and solving of decision-making problems, in accordance with the criterion of *maximum expected utility*. In this way, the ID encapsulates both the Bayesian inference and the decision phase in a single, more flexible and elegant model.

2. Original system

The original system on which the proposed system is based was presented in [8, 9]. This system is a CoRS with argumentative capabilities based on linguistic and cognitive principles. From a design point of view, the original system followed the FANTASIA Interaction model and the PGM of choice were Bayesian Networks (BN), implemented using the aGRuM library [12].

From the knowledge representation point of view, a graph database is adopted to host information derived from Linked Open Data (LOD) sources. For the purposes of this case study, the movies domain will be considered. The knowledge base is constructed by collecting data from different sources and enriched using graph data science techniques, which are employed to capture latent information. The procedure is described in [13]. The main entities of the knowledge base are represented by the labels MOVIE, PERSON and GENRE, which are interconnected by appropriate relationships (such as HAS_GENRE, WORKED_IN, and so on). Additionally, information from the MovieLens 25M² dataset is inte-

²<https://grouplens.org/datasets/movielens/>

grated into the database. A MOVIELENSUSER node is created for each user in the dataset, and a RATED relationship is established between the MOVIELENSUSER node and a MOVIE node for each reported rating in the dataset. In addition to a number of basic properties such as name, year of birth and ratings, MOVIE and PERSON nodes are characterised by authority attributes and hub scores calculated by means of the HITS algorithm [14]. As discussed in [9], these network analysis measures help model cognitive characteristics that are relevant for the selection of *plausible arguments* [15]. Finally, the graph database is used to store a *dialogue state graph* which tracks the agent’s relationships with the knowledge domain and other agents, including humans. This graph can be modified by the agent through speech acts in order to evolve it towards graph patterns that the agent identifies as goal patterns, i.e. a desired configuration of the dialogue state. In this way, the graph database will be interrogated by the CoRS by extracting a relevant sub-graph taking into account the knowledge base and belief of the system evolved during the conversation.

In the reference system, the decision-making level involves a BN dynamically generated on the basis of the extracted relevant sub-graph. In particular, in the case of Movie Recommendation, the actors, films and genres are nodes of the BN, while the (oriented) relations between them represent the causal relations. Initially, each node is initialised by specifying its own CPT, which can either be pre-calculated or derived from parent nodes. This network is used to adjust the exploitation/exploration cycle, typical of recommendation dialogue [16], by taking into account the data extracted from MovieLens (soft evidence) and the feedback gathered through the dialogue with the user (hard evidence). This way, the BN can represent the probability of each movie and each feature to be of interest for a user, after applying Bayesian inference. Based on the information extracted from the Bayesian network, a module outside the PGM is responsible for the decisions taken. Specifically, the system decides whether to recommend a candidate item (*exploitation move*) or, in the case of non-recommendation, to ask the most useful question (*exploration move*), based on the criteria considered in [8]. In the case of exploitation moves, in addition to item recommendation, argumentation is provided based on the three most useful features, whose utility is calculated as the harmonic mean of four (normalised) parameters related to cognitive properties [15].

3. Proposed system

The proposed system based on IDs replicates part of the reference strategy: the aim of this work is to provide a first test of the capabilities of the IDs to handle the problem so we concentrate on the fundamental steps of

Table 1

The capabilities of the original system and the ones replicated in the new system using IDs (in bold).

Tracked beliefs	Question types	Question targets	Scores
Wants	Polar	Movie	Hub
Likes	Open	Actor	Authority
Knows		Genre	Entropy

the original strategy. Table 1 shows the characteristics of the reference system, highlighting the ones reproduced by the proposed system. The approach is inspired by the system presented in [17].

3.1. ID for Movie Recommendation

The current system is based on the previously introduced knowledge base and uses the same principles for the extraction of the relevant sub-graph. The decision-making core of the system is represented by the ID, again dynamically constructed from the relevant sub-graph. In particular, the construction of the ID is divided into two parts. The first part concerns the recommendation branch, along which a decision is made whether or not to make an exploitation move. For each movie, whether a candidate or a secondary film, an uncertainty node is generated, and the same is done for the individuals who are part of those films. In particular, the nodes related to films will be median nodes of the nodes related to the individuals who worked on that film. In addition, the query used to extract the relevant sub-graph returns a collection of votes assigned to films, which is used to apply soft evidence to each of the movie nodes (both target and secondary). For each candidate film, an $EST(Movie)$ uncertainty node related to the estimator operating on that film is contextually generated. Indeed, within an ID it is possible to take into account the *truth* (the best movies in this case) and the estimate on the truth (the estimators on the best movies). Furthermore, the ID also takes into account the uncertainty of the estimator if the CPT of the EST nodes is initialised using the relative confusion matrix. As shown in Fig. 1, this information together will influence the system's decision on a potential exploitation move. Which decision will be made about the *Recommendation* node will depend on the utility function governing the goodness of possible choices. This function defines the utility value of not recommending, i.e. the utility of not performing an exploration move:

$$U_{NoRec} = U_{max} \cdot \frac{1}{1 + e^{nTurn-5}} \in [0, U_{max}] \quad (1)$$

where U_{max} represents the maximum utility that can be given to a choice, while the second contribution is given

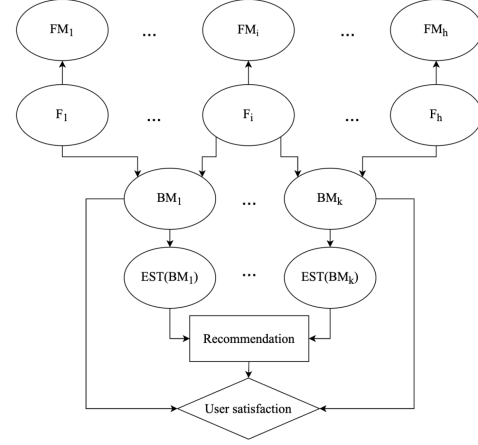


Figure 1: Generic ID structure related to the recommendation branch. *BM* nodes represent best movies, *F* nodes represent features and *FM* nodes represent feature movies, i.e. secondary movies. The topology follows the causal relationships that coexist between the entities involved.

by a sigmoid that takes as input the number of questions asked by the system. The objective is to have a utility of not recommending that is maximum at the beginning of the dialogue and that as the number of questions asked increases, the utility decreases, with an increasing rate of decrease. In this way, the system will be inclined to always ask the user at least one question and never to exceed a certain number of questions. Thus, U_{NoRec} represents the system's indecision with regard to the possible recommendations it can give at that moment, an indecision that is expected to be greatest at the beginning of the dialogue since the system does not yet have any information about the user. In addition to the utility of not recommending, the function defines the utility of recommending a particular candidate movie m :

$$U_m = (2U_{max} \cdot r_m^2) - U_{max} \in [-U_{max}, U_{max}] \quad (2)$$

where r_m represents the rating assigned to the movie candidate m normalised between 0 and 1. In this way, the utility of the recommendation will be linked to the true rating of the candidate movie and its value will be negative for low ratings and positive for high ratings. Thus, recommending an item with a low true rating will be punitive compared to an item with a high true rating. The objective is to prioritise the recommendation of movie candidates with higher true ratings and to disfavour the recommendation of those with low true ratings, possibly by preferring an exploration move.

The second part of the process concerns the exploration branch, during which an exploration move is made. The underlying assumption is that if the utility of "not recommending" is greater than that of recommending

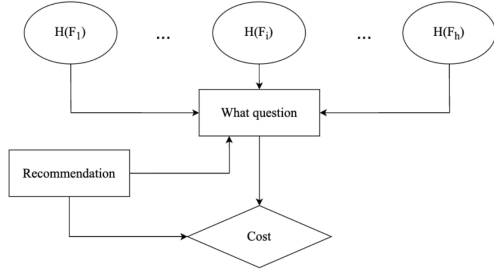


Figure 2: Generic ID structure related to the exploration branch. For each feature (actor, director, etc.), an uncertainty node H is generated, representing its entropy. These nodes, together with the previous decision to recommend or not to recommend, condition the choice of question to ask, which has a cost.

a movie, then a question must be asked. In particular, the most useful question must be chosen. In this case study, as anticipated, the exploration only involves the entropy of the features, not taking into account other aspects of the features and other nodes. In particular, for each feature f extracted from the database, an uncertainty node $H(f)$ is contextually created. Each node H represents the entropy of the related feature. A decision node *What question* is in charge of deciding which question should be asked, and depends on both nodes H and the decision node *Recommendation*, generating a decision sequence starting from the latter. The idea is that the choice of question must depend both on the entropy of the features that can be chosen and on the decision that was made at the time of the recommendation, i.e. the decision to perform an exploitation or an exploration move. Among the possible choices of *What question*, in fact, there is also a *No question*, which only makes sense to choose in the case of an exploitation move. Finally, a *Cost* utility function represents the utility of the *What question* choice. Fig. 2 shows the structure of the exploration branch in a generic form. Tab. 2 shows the cost associated with each decision sequence that the system is capable of undertaking. In particular, the highest cost, equal to -100 , is applied to those decision sequences that are to be avoided. Conversely, the lowest cost, equal to 0 , is applied to the case where the system does not ask questions. A variable cost, between -100 and -1 , is calculated in the case where the system decides not to recommend and ask a question about an actor. The magnitude of this cost will depend on the entropy value of the relevant uncertainty node. The higher the entropy, the lower the cost of the corresponding question. The idea is to collect evidence on the uncertainty nodes on which the model’s uncertainty is most concentrated, as the system’s objective is to lower the model’s entropy level before making a recommendation.

Table 2

The cost function the system considers when deciding to ask questions.

Recommendation	What question	Cost
Movie	Actor	-100
Movie	No question	0
No	Actor	$-99 \cdot (1 - H(f)) - 1$
No	No question	-100

3.2. Simulation

The current system was tested by simulating a dialogue between the system and a MovieLens user whose answers are derived from ratings recorded in the dataset. At the beginning of the conversation, the agent has no information about the user and for this reason the user immediately specifies the preferred genre. This information is derived by searching the database for that genre for which the average rating of that particular user is the highest. All following questions are polar questions and concern PERSON type features. Again, the answer is derived by considering the ratings given by the user to the ARGITEMs associated with that feature. Once the genre is known, a positive belief *likes* is created that associates the user with the preferred genre and at this point the database is queried by extracting the best three, the related features and the secondary films. If, from the ID, the best action is to recommend, the system proposes one of the candidate films to the user; otherwise, if the best action is not to recommend, the system asks the most useful question. If the user’s answer consists of a positive or negative preference, this involves adding evidence in the system, adding the user’s stance on that feature to the database and reconstructing the ID from a dataframe extracted with the same query used at the beginning of the dialogue. The idea is that by keeping track of the user’s stances collected as the system asks questions, it is possible to extract target movies that are more consistent with the user’s preferences. When a film is recommended, the system also provides arguments to support its choice, consisting of a selection of the most important features related to the recommended film, thus implementing Argumentation-based dialogue [18]. The dialogue provided by the simulation is constructed by using templates causing the generated conversation to sound unnatural. For this reason, these template-based dialogues were reformulated by ChatGPT-4 to make the conversation more natural, using the following prompt: *Rephrase the following dialogue to make it sound more natural. Keep the structure and only change the sentences.*

In this task we ask you to evaluate the quality of the conversation between Mary, who is trying to recommend a movie, and George, who is looking for a movie to watch. Read the following dialogues and provide your ratings using the form.

Q1 Is Mary asking coherent questions to help George finding a movie?
 Q2 Are the two people communicating naturally?
 Q3 Does Mary show a good expertise about the movie domain?
 Q4 Assuming he has not seen the movie, do you think George would accept Mary's suggestion?

Figure 3: Survey task and questions posed to participants for each dialogue.

4. Experimental setup

The experimental phase followed the approach used in [8]. The approach involves recruiting 20 participants via the Prolific³ portal who were asked to complete a survey on the Qualtrics⁴ platform that involves the evaluation of 20 dialogues divided into three types:

- Five dialogues taken from INSPIRED Corpus [19], a dataset of human-human interactions for Movie Recommendation. These dialogues represent the positive subset of the control group.
- Five system-generated dialogues where both the extraction of candidate films and the choice of supporting features are random, independent of system belief. These dialogues represent the negative subset of the control group.
- Ten dialogues generated by the system using the presented strategy, which represent the target dialogues.


Fig. 3 shows the survey task with the four questions asked to the participant for each dialogue, for which the participant gives a score between 1 and 5. Q1 refers to the consistency of the questions asked during the exploration move, in order to understand whether the features are selected correctly during the dialogue. Q2 and Q3 refer to the naturalness of the dialogue, with the latter referring to the user's perception of the recommender's level of expertise. Finally, Q4 refers to the quality of the features chosen to support the recommendation. In conclusion, the participants were native English speakers living in the UK or US and they were compensated according to the average hourly wage of their home country.

5. Results


Fig. 4 shows the scores obtained by the current system based on ID for each question blue(b), compared with the scores obtained by the original system based on BN (a). In both instances, the scores obtained by the target

³<https://www.prolific.com/>

⁴<https://www.qualtrics.com/>



(a) Results obtained by the original system



(b) Results obtained by the proposed system

Figure 4: Comparison between the results obtained by the original system based on BN (a) and by the proposed system based on ID (b). The obtained results show higher scores than the baseline represented by the negative dialogues but not as high as the ones obtained by the original system. The difference between the two systems is expected as only part of the original strategy is replicated in this work, excluding a series of significant aspects, such as asking open questions and discussing films as well as the people who work in them.

dialogues are higher than those obtained by the negative dialogues and lower than those obtained by the positive dialogues. In particular, the difference between target and negative dialogues is more pronounced on Q4, which is an indicator that the supporting arguments make the recommendation plausible.

As an objective measure, during the generation of the dialogues for each round, the average normalised entropy of the ID was recorded, calculated as the average of the normalised entropy among all variable nodes of the model. In Fig. 5 it can be observed that a) during a target dialogue the average entropy of the model decreases, in contrast to the case where b) the dialogue is random and the average entropy of the model does not tend to decrease. The first scenario is compatible with the idea that the system accumulates information as the dialogue progresses, in accordance with the strategy adopted. In the second scenario, on the other hand, the ID is regenerated at each turn from randomly extracted candidate films, making it unlikely that the new extracted features contribute in accumulating coherent information.

To further analyse the data concerning the synthetic di-




Figure 5: Trend of normalised mean entropy of the ID during (a) target dialogues and (b) random dialogues. These trends were obtained by measuring the entropy of the system during the generation of ten target dialogues in (a) and ten random dialogues in (b).

alogues, we use a Cumulative Link Mixed Model (CLMM) [20] with Laplace approximation, [21]. This model accommodates random effects attributable to individual participants or specific stimuli, treating them as blocking variables and assesses the likelihood of observing high values on the Likert score in relation to the independent variable (i.e., dialogue type). The test revealed that the association between the occurrence of high scores, in general, is very strong ($p < 0.001$) for both target and positive dialogues and, as expected, absent for negative values. This result is stronger with respect to the results obtained in [8, 9], where only a weak association was observed. There are multiple aspects that contribute to this result, in our opinion. First of all, in the original work, the p -value was already very close to the strong significance threshold ($p = 0.0144$), so the effect was only technically considered *weak* even in that case. Also, there is a chance that the simplified situation may have harmed negative dialogues more than the other two categories. As a final remark, however, the IDs have indeed made the decision process more uniform and flexible, given the introduction of utility functions and a unified framework for decision making. The quality improvement of the decision process management, especially in deciding when to recommend, given the available arguments to support the position has improved the system even in its basic form.

6. Conclusions & future work

The results obtained indicate that the implementation of a knowledge graph exploration strategy based on the ID is more effective than a random strategy. This conclusion is further supported by objective measures, including the system’s entropy, which decreases as the system accumulates information during the dialogue before making a recommendation. It is therefore possible to generalise within an ID a decision-making process that, in the original system, was implemented by a module external to

the probabilistic model. The results achieved in this case were lower than the ones of the original system, but this was expected as only part of the original strategy was replicated. Future work will cover the implementation of the missing functionalities and the deployment of the system in the Unreal Engine, as the technology to implement IDs has been integrated in the FANTASIA plugin. We will also investigate the possibility of integrating the argument selection process in the ID to fully support Argumentation Based Dialogue.

Acknowledgments

This work is supported by the Supporting Patients with Embodied Conversational Interfaces and Argumentative Language (SPECIAL) project, funded by the University of Naples on the “Fondi per la Ricerca di Ateneo” (FRA) program (CUP: E65F22000050001)

References

- [1] A. Darwiche, Human-level intelligence or animal-like abilities?, *Communications of the ACM* 61 (2018) 56–67.
- [2] M. T. Hicks, J. Humphries, J. Slater, Chatgpt is bullshit, *Ethics and Information Technology* 26 (2024) 38.
- [3] J. L. Austin, *How to Do Things with Words.*, Clarendon Press, 1962.
- [4] J. Pearl, D. Mackenzie, *The book of Why*, Basic Books, 2018.
- [5] A. Origlia, F. Cutugno, A. Rodà, P. Cosi, C. Zmarich, Fantasia: a framework for advanced natural tools and applications in social, interactive approaches., *Multimedia Tools and Applications* 78 (2019) 13613–13648.
- [6] G. Flórez-Puga, M. A. Gomez-Martin, P. P. Gomez-Martin, B. Diaz-Agudo, P. A. Gonzalez-Calero, Query-enabled behavior trees, *IEEE Transactions on Computational Intelligence and AI in Games* 1 (2009) 298–308.
- [7] J. Webber, A programmatic introduction to neo4j, in: *Proceedings of the 3rd annual conference on Systems, programming, and applications: software for humanity*, ACM, 2012, pp. 217–218.
- [8] M. Di Bratto, A. Origlia, M. Di Maro, S. Mennella, Linguistics-based dialogue simulations to evaluate argumentative conversational recommender systems, *User Modeling and User-Adapted Interaction* (2024) 1–31.
- [9] M. Di Bratto, A. Origlia, M. Di Maro, S. Mennella, On the use of plausible arguments in explainable conversational ai, in: *Proceedings of Interspeech*, 2024.

- [10] R. A. Howard, J. E. Matheson, Influence diagrams, *Decision Analysis* 2 (2005) 127–143.
- [11] A. Biedermann, F. Taroni, Bayesian Networks and Influence Diagrams, 2023, pp. 271–280.
- [12] G. Ducamp, C. Gonzales, P. Wuillemin, agrum/pyagram : a toolbox to build models and algorithms for probabilistic graphical models in python, in: *Proceedings of the 10th International Conference on Probabilistic Graphical Models*, volume 138, PMLR, 2020, pp. 609–612.
- [13] A. Origlia, M. Di Bratto, M. Di Maro, S. Mennella, A multi-source graph representation of the movie domain for recommendation dialogues analysis, in: *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, 2022, pp. 1297–1306.
- [14] J. M. Kleinberg, Authoritative sources in a hyperlinked environment, *Journal of the ACM (JACM)* 46 (1999) 604–632.
- [15] F. Paglieri, C. Castelfranchi, Revising beliefs through arguments: Bridging the gap between argumentation and belief revision in *mas* (2005) 78–94.
- [16] C. Gao, W. Lei, X. He, M. de Rijke, T.-S. Chua, Advances and challenges in conversational recommender systems: A survey, *arXiv preprint arXiv:2101.09459* (2021).
- [17] P. Lison, C. Kennington, Opendial: A toolkit for developing spoken dialogue systems with probabilistic rules, in: *Proceedings of ACL-2016 system demonstrations*, 2016, pp. 67–72.
- [18] H. Prakken, *Historical overview of formal argumentation*, volume 1, College Publications, 2018.
- [19] S. A. Hayati, D. Kang, Q. Zhu, W. Shi, Z. Yu, Inspired: Toward sociable recommendation dialog systems, in: *2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Association for Computational Linguistics (ACL)*, 2020, pp. 8142–8152.
- [20] A. Agresti, *Categorical data analysis*, volume 792, John Wiley & Sons, 2012.
- [21] Z. Shun, P. McCullagh, Laplace approximation of high dimensional integrals, *Journal of the Royal Statistical Society: Series B (Methodological)* 57 (1995) 749–760.