Culture Awareness in Intelligent Systems

Ariel Gjaci[1], Luca Oneto[1], Carmine Tommaso Recchiuto[1] and Antonio Sgorbissa[1]

[1]*Università degli Studi di Genova, via all'Opera Pia 13, 16145, Genova, Italy*

Abstract

Understanding how the data used to train intelligent systems affects their behaviour is a critical task in the Artificial Intelligence field. It is also known that making artificial agents capable of adapting their actions according to the culture improves their interaction with humans. For this reason, it may be crucial to know how the cultural component inside data affects the prediction of intelligent systems. In this paper, we propose a method to acknowledge the cultural factor inside data, and we show some preliminary results obtained by using Random Forest and Support Vector Machine models on two publicly available datasets.

Keywords

Culture Awareness, Artificial Intelligence, Social Robotics

1. Introduction

Intelligent systems, e.g., social robots, are a technology that has grown in the last few years. According to a recent definition of [1], "an intelligent system operates in an environment with other agents, possesses cognitive capabilities such as perception, action control, deliberative reasoning or language, follows principles of behaviour based on rationality and social norms, and can adapt by learning." A key goal of an intelligent system is, therefore, to interact with other agents, e.g., humans, by exploiting its capabilities.

On the other side, it is well known that people's cultures play a crucial role in human-human interaction: if we consider, for example, the thumb-up gesture, we notice that this is a widely recognized sign of approval in several countries, but it is also used as an insulting gesture in Bangladesh. Hence, given the importance of culture in *human-human interactions*, it is natural to ask oneself if culture may play the same important role also in *human-robot interactions*. Answers to this question have been given in some recent works [2], [3], which have assessed how culture influences people's expectations, attitudes, and behaviours before, during, and after the interaction with robots. Overall, it appears that a social robot that displays cultural sensitivity promotes human acceptance.

For this reason, the problem of making artificial agents able to interact with humans by considering their culture has also been recently addressed. To embed culture in intelligent systems, researchers might use two approaches:

9th Italian Workshop on Artificial Intelligence and Robotics (AIRO 2022)

ariel.gjaci@edu.unige.it (A. Gjaci); luca.oneto@unige.it (L. Oneto); carmine.recchiuto@dibris.unige.it (C. T. Recchiuto); antonio.sgorbissa@unige.it (A. Sgorbissa)

0000-0002-8445-395X (L. Oneto); 0000-0001-9550-3740 (C. T. Recchiuto); 0000-0001-7789-4311 (A. Sgorbissa)

- Theory-based: design a cultural representation that is embedded inside computational models such that they will produce culturally-aware predictions.
- Data-driven: exploit culturally-dependent data in the learning phase of models so that they can produce culturally-aware predictions.

Among the most relevant works in this context, a theory-based model for automatically adapting HCI systems to users by using a culturally-aware adaptive procedure has been presented in [4], while Papadopoulos et al. [5] have built a theory-based system that can be integrated into social robots for applying an understanding of the culture, customs and etiquette of the person who is interacting with the robot while autonomously reconfiguring the way of acting and speaking. Speaking about data-driven approaches, Lugrin has combined the theory-based approach with the data-driven one to create a culturally-dependent Bayesian Model for virtual characters [6], while in [7] a data-driven algorithm was developed to generate culture-aware co-speech gestures in social robots. However, all the aforementioned works do not provide any detail on how much the cultural component affects the data. This information can be important to transfer the knowledge acquired by a model trained with culture-dependent data to another culture [8], or to make the model fair with respect to the cultural component with the recent fairness techniques [9]. Addressing this problem can be useful, especially when the availability of culturally-dependent data is limited.

Following the example at the beginning of this article, we may consider a classification problem with a culturally-dependent gestures dataset where the goal is to predict whether the gestures are offensive. By analysis, we may find that the accuracy of a model trained with data depending on one culture and tested with data of the same culture is significantly higher than the accuracy of the same model tested with data depending on another culture. However, building an accurate model for every different culture may not be possible since the labeled data are insufficient. For this reason, we may exploit an instance weighting technique [8] to transfer the knowledge of the first model to the second one by also considering the cultural difference found previously. Please note that if we instead find that the classification difference is not significant between the two test data, we may also use the same model to classify both cultures.

By combining the two concepts of intelligent systems and culture, we propose in this work a methodology based on classification techniques for understanding the role of culture when intelligent systems need to be trained to perform a given task, for example, activity recognition. Our method starts from the conjecture that, in a classification problem, the data to be classified are necessarily affected by cultural aspects. This method can be, in principle, used with any dataset containing a feature that clearly identifies the cultural context where data have been collected.

In the following chapter, we show some preliminary results obtained by relying on custom Random Forest and Support Vector Machines models with two different culturally-dependent datasets.

2. Methodology

The proposed work aims to analyze how the cultural component affects the data by evaluating a learning model trained on different datasets. More in detail, we picked two public datasets [10] and [11]. In both datasets, we spot a feature identifying the culture x_c inside the data sample $x \in \mathcal{X} \subseteq \mathbb{R}^d$ used to predict output $y \in \mathcal{Y} \subseteq \{\pm 1\}$, where d is the number of features. During each test, we focused on a subset of the possible cultures to perform pairwise comparisons (e.g., Italian vs. Chinese or Jamaican vs. German), and therefore we extracted a set of samples and features for each dataset to be considered for classification. Specifically, we started choosing two admissible values a,b for the cultural identifier x_c, and we considered only the samples for which $x_c \in \{a, b\}$.

At the end of this process, we considered the following datasets:

- The CROCUFID food dataset [10] has the area of origin of the food as the cultural identifier with values $\{Asian, Western\}$, the domain $\mathcal{X} = \{$*22 visual features that describe the food*$\}$, and the prediction $\mathcal{Y} = \{$*the food is sweet or savoury*$\}$. That is, starting from the visual appearance of a food, we aim to predict if the food will be sweet or savoury, and we expect this depends on cultural aspects.
- The whats-cooking recipes dataset [11] has the type of the cuisine $\{Italian, Cajun - Creole, Thai, Japanese, etc.\}$ as the cultural identifier, the domain $\mathcal{X} = \{$*58 ingredients that may be present/absent in the dish*$\}$, $\mathcal{Y} = \{$*presence/absence of a critical ingredient for the person's health, e.g., sugar, butter, ingredient to which the person is allergic*$\}$. That is, starting from known ingredients, we aim to predict if a critical ingredient may be present or not, and we expect this depends on cultural aspects.

Both datasets would be helpful to predict whether or not a food contains a critical ingredient for a person's health. For instance, a social robot that suggests dishes to a person will avoid suggesting those that contain sugar if the person is overweight or diabetic.

We trained and optimized four Random Forest and Support Vector Machines models:

- **1)** Trained by using only the data that have $x_c = a$ and by removing the cultural identifier.
- **2)** Trained by using only the data that have $x_c = b$ and by removing the cultural identifier.
- **3)** Trained by using all the data and by removing the cultural identifier.
- **4)** Trained by using all the data.

To choose the best Random Forest model, we used 1000 trees, and we optimized only the hyperparameter that controls the cardinality of the random subset of features that each node of the trees can check. To choose the best Support Vector Machine model, we used the radial kernel type, and we optimized the *C* hyperparameter that controls how much we want to avoid misclassifying each training example and the *gamma* hyperparameter that controls the kernel size. We used 10-fold cross-validation and Balanced Accuracy as well as AUC value as metrics for both models.

After the training, we tested each model by using two different test sets: one contains only data that have $x_c = a$ while the other contains only data that have $x_c = b$. To better compare the models, we maintained constant the number of training and validation samples that we used.

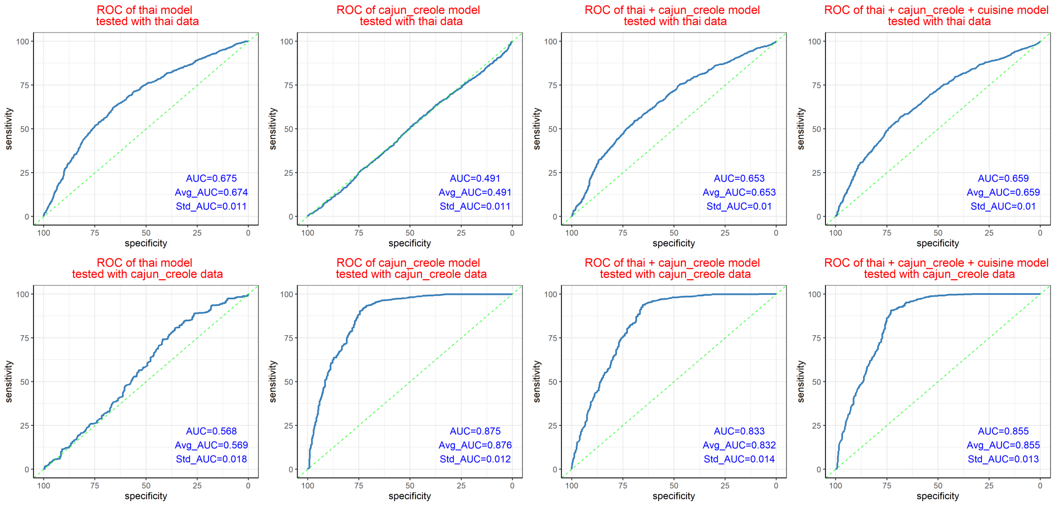

Figure 1: ROC plots of the whats-cooking dataset trained and tested with data coming from Thai and Cajun-creole cuisines by using Random Forest models. The prediction is whether the sugar is present in the recipe. In each plot, it may be observed the Area Under the Curve value as well as its average and standard deviation values computed on 500 bootstrap data samples.

3. Results and Future Work

The ROC curves with AUC values achieved by predicting if there is sugar or not in the whats-cooking dataset are shown in Figure 1: it can be noticed that models trained with data related to a specific culture fail to predict the right class if tested with the data related to the other culture. It can also be noticed that adding the *cuisine* feature in the model trained with all the data slightly improves the result.

These results are coherent with the confusion matrices in Figure 2 obtained by optimizing the Random Forest models as described before. The only difference is that in this case, the metric chosen for the optimization is the Balanced Accuracy, which considers possible unbalances of classes during the training. It may be noticed from Figure 2 that the model trained with Thai cuisine data well predicts test data related to Thai cuisine, but does not perform well with the test data related to Cajun-Creole cuisine. The same also occurs in the opposite situation (i.e., model trained with Cajun-Creole cuisine data and tested with data related to Thai cuisine). Finally, the model trained with data related to both cuisine types performs well with both test data, but here an interesting fact can be noticed: the model fails at predicting if the sugar is not present in the test data related to Thai cuisine. This may be due to the fact that both datasets are used in the training process, and the model cannot completely identify them in the classification process.

By using the CROCUFID dataset with Random Forest models, we have obtained results similar to the ones represented in Figure 1 but less noticeable: the model trained with Western data had only an improvement of 0.13 in the AUC value when tested with Western data with respect to Asian. Similarly, the model trained with Asian data had only an improvement of 0.10 in the AUC value. By using the Support Vector Machine models, as shown in Figure 3, the differences

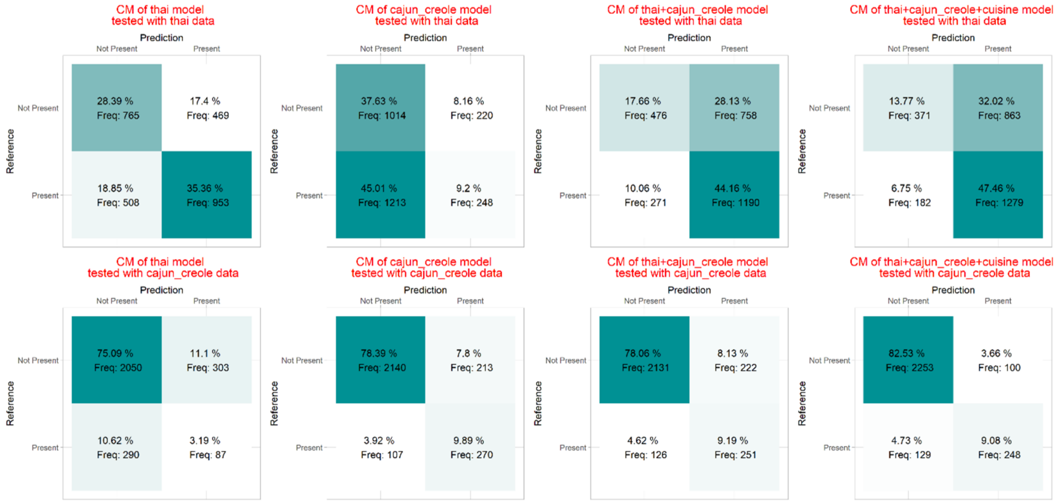

Figure 2: Confusion matrices of the four Random Forest models obtained by using data from Cajun-Creole and Thai cuisines. The prediction is whether the sugar is present in the recipe. Freq indicates the frequency values.

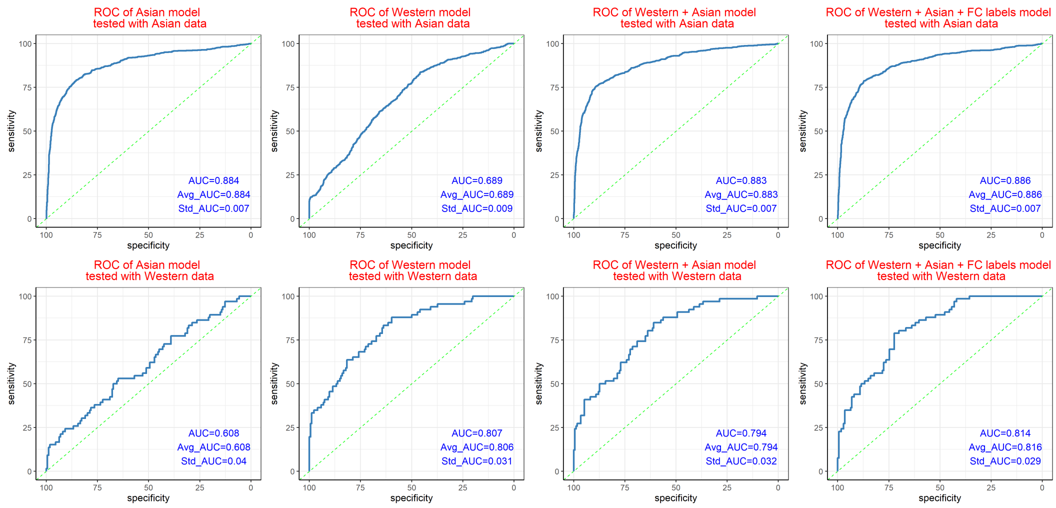

Figure 3: ROC plots of the CROCUFID dataset trained and tested with data related to Asian and Western foods by using Support Vector Machines models. The prediction, in this case, is whether the food is Sweet or Savoury. The curves in the second row are less smooth because the test set containing Western foods has fewer samples. Here we identified a cultural dependency in the feature related to Food Classification, e.g. Asian or Western.

become more pronounced, but the behaviour remains coherent with the one obtained using Random Forest models.

Even if these preliminary results are encouraging, it must be noticed that considering different couples of cultures may produce different outcomes. Sometimes differences between models cannot be noticed, e.g., in the comparison of Indian vs. Italian recipes. The same occurs with other critical predicted ingredients, such as peanuts. Further analyses are needed to better assess the role of cultural factors inside data and their possible exploitation in classification problems. To this aim, a custom dataset related to human gestures has been developed, and it is currently under analysis. The dataset is composed of more than 2000 images, related to two gestures that are similar to each other: the *Namaste* gesture, which is used in the Indian culture, and the *Pray* gesture, typical in different countries around the world.

Overall, this preliminary work lays the basis for a deeper investigation to understand the role of culture in classification problems. We finally note that the authors who created the culturally-dependent datasets may have introduced their stereotypes during the creation process: Italian food is not the one you find in Italian restaurants around the world, and the same is true for other cuisines. We did not consider the possibility that datasets have been compiled with the version "for foreigners" of popular dishes, but this can be explored in the future.

References

[1] M. Molina, What is an intelligent system?, arXiv preprint arXiv:2009.09083 (2020).

[2] M. Salem, M. Ziadee, M. Sakr, Marhaba, how may i help you? effects of politeness and culture on robot acceptance and anthropomorphization, in: 2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI), IEEE, 2014, pp. 74–81.

[3] V. Lim, M. Rooksby, E. S. Cross, Social robots on a global stage: establishing a role for culture during human–robot interaction, International Journal of Social Robotics 13 (2021) 1307–1333.

[4] R. Heimgärtner, Culturally-aware hci systems, in: Advances in culturally-aware intelligent systems and in cross-cultural psychological studies, Springer, 2018, pp. 11–37.

[5] C. Papadopoulos, N. Castro, A. Nigath, R. Davidson, N. Faulkes, R. Menicatti, A. A. Khaliq, C. Recchiuto, L. Battistuzzi, G. Randhawa, et al., The caresses randomised controlled trial: exploring the health-related impact of culturally competent artificial intelligence embedded into socially assistive robots and tested in older adult care homes, International Journal of Social Robotics 14 (2022) 245–256.

[6] B. Lugrin, J. Frommel, E. André, Combining a data-driven and a theory-based approach to generate culture-dependent behaviours for virtual characters, in: Advances in Culturally-Aware Intelligent Systems and in Cross-Cultural Psychological Studies, Springer, 2018, pp. 111–142.

[7] A. Gjaci, C. T. Recchiuto, A. Sgorbissa, Towards culture-aware co-speech gestures for social robots, International Journal of Social Robotics (2022) 1–14.

[8] K. Weiss, T. M. Khoshgoftaar, D. Wang, A survey of transfer learning, Journal of Big Data (2016).

[9] L. Oneto, S. Chiappa, Fairness in machine learning, in: Recent Trends in Learning From Data, Springer, 2020, pp. 155–196.

[10] A. Toet, D. Kaneko, I. De Kruijf, S. Ushiama, M. G. Van Schaik, A.-M. Brouwer, V. Kallen, J. B. Van Erp, Crocufid: A cross-cultural food image database for research on food elicited affective responses, Frontiers in psychology 10 (2019) 58.

[11] whats-cooking, https://www.kaggle.com/c/whats-cooking, 2022. Accessed: 2022-10-04.