

# **Applying the ALTAI framework to a credit scoring algorithm for mortgages - A case study with the Dutch Volksbank**

Cristian Barru , Atia Cort s, Francesca Foffano, Christoph Heitz, Joris Krijger, Long Pham  
Teresa Scantamburlo, Evert Stamhuis

## **Introduction**

The current research assesses the usability of the ALTAI framework for the credit risk scoring model as it is used by the Dutch bank ‘De Volksbank’. De Volksbank is a Dutch state-owned bank owning several bank brands within the Netherlands. It is the fourth-largest bank in the Netherlands, having 3.2 million customers and nearly 3,000 employees. De Volksbank provides mortgages (  47.8 billion in 2018), manages savings (  37 billion) and offers 1.5 million customers a current account.

One of the key risk factors in handling mortgages is credit risk. This is the risk of customers going into default, which brings major inconvenience to the customer and additional and undesired handling costs for the bank that may eventually lead to foreclosure, involving even more substantial losses. Therefore, an important part of the mortgage application procedure of every bank is a Credit Approval or Credit Risk Scoring model. This model is one of the gatekeepers of the application process, influencing which customers should be accepted and which customers should not. The purpose of the credit risk model is twofold: it assists acceptors in achieving a target mortgage portfolio but also assists in regulatory capital calculations. Based on several parameters (such as previous arrears, expenses over income and Loan to Value) a score is retrieved which provides an estimate for the probability of default in a specific period after the loan becomes active. The score indicates to the bank experts but is never in and of itself grounds for the acceptance or rejection of an application. The final decision is based on internal policy guidelines and a separate assessment of an acceptor.

## **Methodology**

In the case study, we investigated two different perspectives: The ‘assessment perspective’ targets the outcome of the Trustworthy AI assessment concerning the investigated algorithm of Volksbank: what could be learned about the trustworthiness of the credit scoring algorithm in the concrete setting? In contrast, the ‘methodological perspective’ focuses on the usability of the chosen assessment method ALTAI.

## **Results**

Assessment: In terms of the trustworthiness of the credit scoring algorithm and ethical aspects of its usage, the assessment led to the following main findings:

- Explainability:
  - Model output, in the form of a red, amber or green traffic light, is used as only one element for a final loan decision, taken by a human decision maker. Applications can be approved even for a red light, or denied even in case of a green light. The final responsibility for the decision is taken by a human decider. The explainability of the final decision is thus only loosely coupled with the explainability of the algorithm.

- Full transparency towards users, i.e. loan acceptors, was deemed undesirable given the risk of gaming, which would compromise the accuracy of the credit risk estimation. The goal of the bank-internal communication about the algorithm is to (a) explain the principal mechanisms, but (b) hide the details, and how exactly the input data influences the output.
- There is no systematic training of users for how to evaluate or reverse engineer decisions by the credit approval model use the system. Only supervisors of users have been informed during so-called “road-shows”, giving the freedom of whether and how much to inform their groups about the inputs of the algorithm. Also, no systematic approach for communication with users about their experience with the model and the output of the model was found in the interviews.
- Fairness
  - In two instances algorithmic bias was signalled and reported. Resolution of these biases was not sought in alterations to the algorithm but in specific user guidelines. These cases also did not trigger a systematic search for algorithmic bias at established internal oversight agencies, such as the model governance department.
  - With respect to algorithmic fairness, it became clear that there is a high level of awareness of this issue, and that it is accepted that unfairness should be avoided. However, it also turned out that there is no consolidated definition of what fairness means concretely, and what kind of unfairness should be avoided. Despite an elaborated reporting system that allows measuring the system’s ability to discriminate between high and low credit worthiness no metrics for fairness are added to the system.

Usability of ALTAI: During the interviews, some difficulties of the ALTAI assessment became apparent:

- ALTAI does not specify to whom the questions should be addressed within the organization. One has to determine beforehand which roles are relevant to interrogate.
- For the ALTAI questionnaire to fit with the specific use-case in the context of finance some preprocessing was required in the order, wording and relevance of the questions. This relates to the first point: when organizations have separate model validation teams, cybersecurity teams and developer teams parts of the assessment (e.g. on robustness, security and data usage) should be done with members from different teams. Finding the match between the assessment and the industry (with industry standards and practices) combined with the expertise of the interviewee is, therefore, a challenge in and of itself.
- The questions are closed (yes/no) and formulated in a very generic way. This required elaboration and translation on the part of the interviewer in order to have a meaningful conversation with the interviewee. Still, confusion arose around terms such as risk ( interpreted as “credit risk” by the interviewees, as opposed to “risk for humans, created from the application of the algorithm”) and “discrimination” (understood as the usual ML interpretation, i.e. referring to differentiating between different groups (good and bad debtors), instead of referring to the systematic disadvantage of socially salient groups). Furthermore, after remarks from an

interviewee that the interview already deepened awareness on ethical questions, we attribute much of this increased awareness to the assessment form (an interview, dialogue) and believe that a self-assessment questionnaire would have much less impact. This leads us to conclude that usability of ALTAI as an independent self-assessment is limited but that it could provide valuable input for ethical evaluations via interviews and interrogations.