JOURNAL OF
CROSS-DISCIPLINARY RESEARCH
IN COMPUTATIONAL LAW [CRCL]

Artificial intelligence and renegotiation of commercial lease contracts affected by pandemic-related contingencies from Covid-19. The project A.I.A.Co.

Maurizio Parton* Marco Angelone* Carlo Metta** Stefania D'Ovidio*
Roberta Massarelli* Luca Moscardelli* Gianluca Amato*
Cristiano De Nobili***

Abstract

This paper aims to investigate the possibility of using Artificial Intelligence (AI) to address the legal issues that arose from the Covid-19 emergency: contracts with continuous, repeated, or deferred performance and, more generally, having to deal with exceptional events and contingencies. We first examine whether the Italian legal system permits 'maintenance' remedies to address contingencies and to avoid the termination of fixed-term contracts whilst continuing to ensure effective protection for both parties' interests. We then provide a complete and technical description of an AI-based predictive framework, designed to assist both the Magistrate (during the trial) and the parties themselves (in out-of-court proceedings) in the redetermination of the rent of commercial lease contracts. This framework, named A.I.A.Co. for Artificial Intelligence for contract law Against Covid-19, was developed under the Italian public grant known as *Fondo Integrativo Speciale per la Ricerca* and – even though the predictive system was initially intended to deal with the very specific problem connected to Covid-19 – the acquired knowledge, the produced model and the research outcomes can be readily transferred to other civil issues (such as, for example, determining the amount of maintenance or divorce obligation in family law).

Keywords: Artificial Intelligence; equitative algorithms; commercial lease contracts; predictive justice; computational law.

Replier: Liane Huttner • Université Paris-Saclay, huttner.liane@gmail.com .

Journal of Cross-disciplinary Research in Computational Law

DOI: pending

www.journalcrcl.org

* University of Chieti-Pescara. maurizio.parton@unich.it, marco.angelone@unich.it, stefania.dovidio@unich.it roberta.massarelli@unich.it, luca.moscardelli@unich.it, gianluca.amato@unich.it.
** ISTI-CNR Pisa, carlo.metta@isti.cnr.it.
*** Pi School. cristiano.denobili@gmail.com.

The A.I.A.Co. project

This paper aims to investigate the possibility of using Artificial Intelligence (AI) to address the legal issues that arose from the Covid-19 emergency: contracts with continuous, repeated, or deferred performance and, more generally, having to deal with exceptional events and contingencies. We first examine whether the Italian legal system permits 'maintenance' remedies[1] to address contingencies and to avoid the termination of fixed-term contracts whilst continuing to ensure effective protection for both parties' interests. We then provide a complete and technical description of an AI-based predictive framework, designed to assist both the Magistrate (during the trial) and the parties themselves (in out-of-court proceedings) in the redetermination of the rent of commercial lease contracts. Even though the predictive system was initially intended to deal with the very specific problem connected to Covid-19, the acquired knowledge, the produced model and the research outcomes can be readily transferred to other civil issues (see Section 'Limitations and future developments' for further details).

The Italian legal framework

The pandemic and the lockdown-type containment measures implemented by the Authority to prevent infections, as extraordinary, unforeseen, and unforeseeable events, can, in fact, be qualified as 'contingencies' (*force majeure* and *factum principis*). Consequently, jurists have had to deal, on the one hand, with the distribution of the contractual risk and, on the other hand, with the management of contingencies from a perspective of preservation of the contract. On this point, see Benedetti [2020], Dolmetta [2020a], Federico [2020], Grondona [2020], and Macario [2020a,b].

The Italian Civil Code provides remedies to address both contingencies and contract breaches.

Regarding the first profile, the primary focus should be on the Termination for Supervening Impossibility (art. 1256 and 1463 of the civil code and following) and the Termination for Excessive Burden (art. 1467 of the civil code). Similar 'demolition' remedies are consequences of the non-fulfillment of the obligation: in such cases, attention should be directed towards the debtor's liability (art. 1218 of the civil code and following) and to the Termination for Breach of Contract under art. 1453 of the civil code. These remedies only grant the possibility to terminate the contractual relationship to the disadvantaged party. However, contract termination does not always best serve the interests of the parties involved, who may prefer to continue the contractual relationship with modified provisions. Therefore, it is essential to examine the existence of 'maintenance' remedies that allow for addressing contingencies and avoiding contract termination while safeguarding the interests of both parties.

In this context, Italian legal scholars have debated the possibility of imposing a duty to renegotiate on the parties, even in the absence of an explicit normative provision (for an overview of different viewpoints, refer to Prisco [2021]). This duty already exists in other civil law systems, such as the German Civil Code (par. 313 BGB about the so-called *Störung der Geschäftsgrundlage*) and the French Civil Code (art. 1195 which is about the so-called *imprévision*).

According to some Italian jurists, potential solutions aimed at preserving the contract can already be found within the Italian Civil Code, see Macario [1996] and Perlingieri [2017, 2011, 2020a]; *contra* Gentili [2003] and Sicchiero [2002]. The obligation to renegotiate, in particular, could be inferred from the general clauses of good faith and fairness pursuant to the art. 1175 and 1375 of the Italian Civil Code and from the principles of proportionality, reasonableness, and adequacy, as well as the principle of 'social solidarity' pursuant to art. 2 of the Italian Constitution, which has been recognized as directly applicable in private relationships [Lipari 2003; Perlingieri 2001, 2020b]. Therefore,

[1] In order to deal with contractual defects or supervening events that affect the validity or equilibrium of the contract, Italian contract law provides for so-called 'demolition' remedies (such as nullity, invalidity, rescission and termination), which aim to remove the effects of the contract itself. However, in addition to these remedies, interpreters have recognised the existence of 'maintenance' remedies, which aim to avoid the termination of the contract and to ensure its preservation. These include the obligation to renegotiate the contract, which can be derived from the general clauses of good faith and fairness *ex* art. 1175 and 1375 of the Civil Code, as well as from the principles of proportionality, reasonableness and adequacy, and the principle of solidarity enshrined in art. 2 of the Italian Constitution, which must be recognised as directly applicable in relations between private parties.

from this perspective, the static and formalistic principle of '*pacta sunt servanda*', codified in the art. 1372 of the Italian Civil Code, should be abandoned in favor of the different principle of '*rebus sic stantibus*'.

In addition, the proposed reform of the Italian Civil Code [Conte 2019] already granted the parties the right to request contract renegotiation, even before a Court, when one party's performance becomes excessively burdensome [Sirena 2020]. The Italian Legislator also intervened on this point during the Covid-19 emergency, albeit only in some specific cases, introducing an obligation to renegotiate[2].

In the international context, regulations governing the performance of cross-border contracts emphasize the implementation of maintenance remedies through special 'hardship' clauses that require parties to renegotiate contracts in the event of contingencies[3]. Furthermore, in 2003, the International Chamber of Commerce developed specific '*force majeure*' and 'hardship' clauses, which were updated in 2020 due to the pandemic [International Chamber of Commerce 2020]. Consistently, the Principles for the Covid-19 crisis (n. 13), formulated by the European Law Institute [2020] during the health crisis, recommend that States ensure renegotiation between parties, even in the absence of a specific contractual clause or legislative provision, in accordance with the principle of good faith.

Methods and results

What and why The A.I.A.Co. (Artificial Intelligence for contract law Against Covid-19) project originated from a competitive research grant called '*Fondo Integrativo Speciale per la Ricerca*' - FISR 2020 (Project Code: FISR2020IP 04568; CUP: D25F21000500007) – provided by the Italian Ministry of University and Research to address the legal issues arising from the Covid-19 emergency. On April 30, 2021, A.I.A.Co. passed the initial selection phase, leading to the development of the prototype described in this paper. In May 2024, the Evaluation Committee gave a favourable opinion on the activities carried out in the framework of the presented project, thus paving the way for the possible funding of the second phase of the project.

A.I.A.Co.'s final aim is to establish a system of 'equitable algorithms' for commercial lease contracts, that is, an AI-based framework that can facilitate the redetermination of rents, both in judicial and extrajudicial contexts, while ensuring the preservation of contracts that have been significantly impacted by the health emergency.

We restrict to commercial (real estate) lease contracts because this is one of the main topics of the jurisprudential debate that has emerged in Italy following the pandemic [Carapezza Figlia 2020; Dolmetta 2020b; Pisu 2021]. By limiting several business activities, the lockdown-type containment measures placed tenants in financial difficulties when it came to paying rent. In response, landlords began enforcing the guarantees that tenants had signed to cover payments, seeking contract terminations, or initiating eviction procedures for delayed rent payments. In many cases, judges, striving to balance conflicting interests, have deemed it equitable to reduce the rental fees, attributing the case, alternatively, to the supervening (partial or temporary) impossibility to perform by the owner[4] or to the obligation to renegotiate the contractual services, inferred from the principles of good faith and fairness during contract execution[5].

[2] See the art. 216, paragraph 3, d.l. 19/10/2020, n. 34, conv. l. 17/07/2020, n. 77 regarding the reduction of rental fees for sports facilities during the lockdown period; the art. 6-*novies* d.l. 22/03/2021, n. 41, conv. l. 21/10/2021, n. 69, addressing the right to renegotiate commercial lease contracts for businesses significantly impacted by the pandemic; the art. 10 d.l. 24/08/2021 n. 118, conv. l. 21/10/2021, n. 147, concerning the negotiated composition of business crises.

[3] This is also recognized within the *Unidroit* Principles (art. 6.2.3), the Vienna Convention of 1980 (art. 79), the Principles of European Contract Law (art. 6:111), the Common Frame of reference (art. 108-110), the European Code of Contracts (art. 157), and the Draft United Nations Code of Conduct on Transnational Corporations (art. 11).

[4] On which see, *ex multis*, Court of Roma, 29/05/2020; Court of Roma, 25/07/2020; Court of Venezia, 28/07/2020; Court of Venezia, 30/09/2020; Court of Milano, 28/06/2021, n. 4651; Court of Reggio Calabria, 29/03/2022, n. 373; *contra* Court of Roma, 09/09/2020.

[5] On which, see, *inter alia*, Court of Milano, 08/04/2020, n. 2319; Court of Roma, 27/08/2020; Court of Milano, 21/10/2020; Court of Auditors, Emilia-Romagna, 03/06/2021; Court of Palermo, 09/06/2021; Court of Lecce, 24/06/2021; *contra ex multis* Court of Roma, 21/05/2021, n. 9457 and Court of Roma, 06/08/2021.

Who and how. The research team was divided into a Law Unit and an AI Unit, which worked collaboratively and continuously shared the progressively achieved results. This organizational structure has proven to be highly effective in terms of project management, particularly considering the interdisciplinary nature of the analyzed issues, which necessitated a combination and integration of various skills and expertise.

Ideally, the A.I.A.Co. framework should take the defense deeds as input and output a decision to be proposed to the judge, along with explanations of the proposed decision. The decision could be a numerical value for quantitative legal prediction problems or a binary or multiclass decision. In any case, such a framework should start with a preprocessing phase of the defense deeds, followed by a natural language processing (NLP) phase extracting relevant features, followed by a predictive regression or classification model. One might question why it is necessary to use such a complicated design instead of an easy-to-implement predictive model that uses a spreadsheet. Well, because the latter predictive model would not be learned: it would require experts to choose (*ex ante*) all the different parameters and scenarios and the underlying rules.

The predictive model should then be analyzed by explainability techniques producing explanations. See Figure 1 for details.

The two research units initially decided to solely focus on the final prediction at first, because NLP is notoriously very expensive in terms of time and resources when fully trained from scratch. We thus decided to postpone the implementation of an NLP phase and to adopt, for the prediction phase, training data artificially generated and submitted for labeling to authentic judges. We point out that, since in this prototype we are interested in the redetermination of the rents, which are numerical quantities, we need to frame the problem as a regression task.

For explanations, we used SHAP. In Figure 1, the prototype described in this paper is the non-greyed-out part of the picture. In the next paragraphs, we provide more detail on how the training data was produced, which predictive model we opted for and how the explainability part works.

Generation of training data. The process of generating our training data was meticulous and time-consuming. This took approximately six months, which reflects the importance we placed on producing representative and meaningful data. We began with an expert-based selection of the most informative features that the judge usually has to evaluate in his decision, identified in 25 fields, see Figure 6 in annex. For each feature we then selected a corresponding marginal distribution. This features and marginals selection process, carried out by the Law Unit[6] in consultation with a small group of judges and the AI Unit, entailed a thorough study and analysis of decisions that had already been published. It was a lengthy process, but it was necessary to ensure the relevance and accuracy of our data.

During these six months, there was continuous interaction between the Law Unit, the judges, and the AI unit. This collaboration was crucial to be able to refine the features and to ensure that they were representative of the problem at hand.

In addition, in the final survey (Section 'Validation survey'), we sought further validation of these features by ask-

[6] This Unit is composed of scholars and PhD students in legal disciplines.

[7] Below one of the artificially generated cases. The parts in italics identify the selected features, that randomly differentiate each case. 'Case
WHEREAS:
- Mr. Marco Rossi, *a private subject*, is the tenant of a commercial premises in Bari, in which he carries out the activity of *Management of multipurpose sports facilities* (ATECO code *93.1*), *the only source of his income*;
- the lessor is *Ente Alfa*, a *private subject*;
- the contract provides for a monthly rental fee of *5.600,00 euros*, to be paid in *monthly* installments of *5.600,00 euros* each;
- that rental fee is the *65%* of the lessor's total annual income;
- the contract *provides the tenant's right to withdraw* (in addition to the prevision of the Article 27, Paragraph 8, Law No./ 392 of 1978), and *does not provide for an express termination clause, in case of not-fulfilled payment of rent*;
- the contract, for the fulfillment of the obligations of the tenant, provides for *a guarantee given by a non-professional party*;
- during the lockdown from Covid-19, *the tenant's income was reduced in 9 months by 85%*;
- the tenant *did not benefit from any income support measures*;

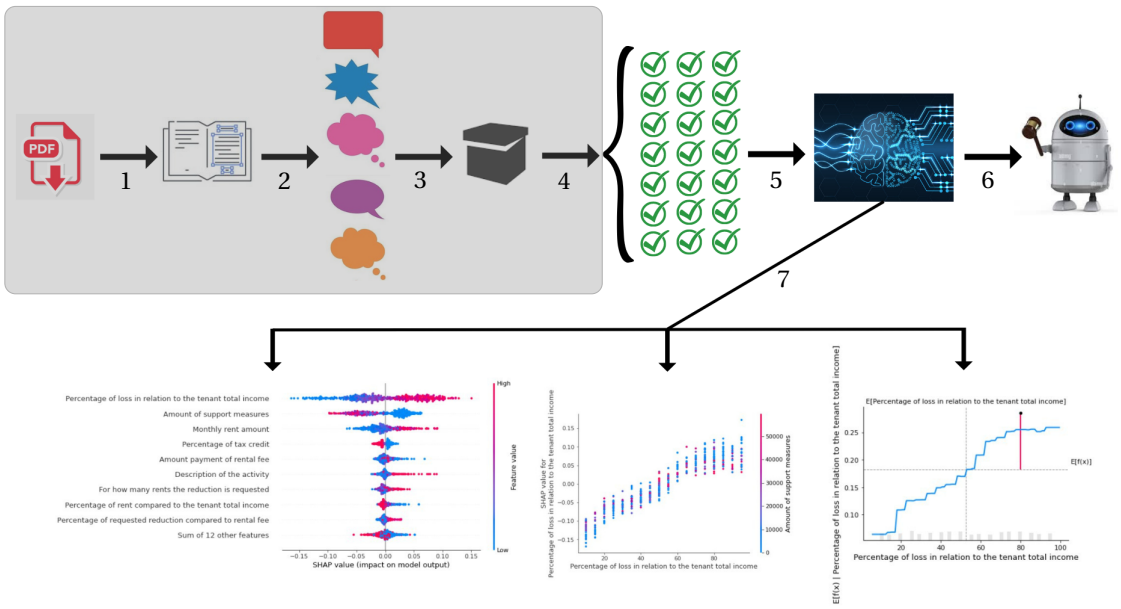

Figure 1: *The complete A.I.A.Co. framework. In 1, a document containing the defense deeds is preprocessed, producing a text (2) that is then passed to a NLP model (3). In 4, the NLP model outputs the relevant features, which are then passed (5) to the predictive model, which outputs (6) the decision that is then proposed to the judge. Weights learned during the training of the predictive model are passed (7) to SHAP to produce explanations helping the judge in using the proposed decision.*

ing judges, lawyers, scholars, and students in legal disciplines to confirm them[7]. We stress that the survey does not address the distribution shift problem, because the participants see only one sample. The distribution shift problem is likely mitigated by the protocol described above.

At the end, we obtained a set of realistic constraints to generate likely cases, such as, for example: 'monthly rental fee amount' greater than € 500 and less than € 50,000, or: if 'nature of the owner' = 'natural person' then 'quality of the owner' = 'private". This protocol was used to generate 600 data points, that is, 600 points in $\mathbb{R}^{25}$, sampled from features distributions with weights chosen again by the Law Unit, the AI Unit and the judges. We expect this human-in-the-loop approach to massively reduce the inevitable distribution shift from the training data to the distribution of actual cases seen by judges.

Each of these data points was used to produce a likely, artificially generated case, similar to a defense deed in a human-readable PDF document. We then, again followed a human-in-the-loop approach by sending these docu-

- the tenant has obtained - for the period indicated in the law - *a tax credit of 60% of the paid rental fees*;

- in the absence of an agreement between the parties on the renegotiation of the contract, the tenant requests to this Judicial Authority the reduction, according to equity, of the amount of the *monthly rental fee for 9 months*.

FOR THESE REASONS

The Judge [...]'.

At that point there was a 'drop-down menu' with the following alternatives: '*DOES NOT ORDER the reduction of the monthly rental fee*'; or '*ORDERS the reduction of the [...] (with the possibility of indicating a percentage between 5% and 100%) of the monthly rental fee*'.

ments to a selected panel of 30 judges[8] for labelling with a percentage representing the amount of the reduction. In order to obtain more accurate labels, we opted to solely include professional judges in the panels, almost all of whom were active in the ordinary jurisdiction (96.67%). At the end, we received 557 answers. Subsequently, we used this labeled dataset for training and testing. We acknowledge that human decisions can be influenced by bias and it is not uncommon for different judges to render dissimilar rulings on the same issue, so there is a risk that automated decisions could reproduce such biases. In order to mitigate both human and automated biases, the proposed tool seeks to integrate both human and machine assessments (see also section 'Impact to justice').

Models, hyperparameters and training. The aforementioned dataset of 557 labeled data points was used to train and test several baselines: a 'constant' model that always predicts the same percentage of reduction, a model that predicts the median reduction on the training set, a decision tree, and a linear model.

We then trained a random forest, with 10-fold cross-validation. This initial training was used to identify and dismiss the less important features. For that purpose, we used a threshold of 10^{-5} for SHAP values, resulting in a new dataset with 21 features instead of 25. See Figure 6 of the annex and the explainability paragraph in this section.

Finally, we trained and tested a random forest and a neural network, again with a 10-fold cross-validation as train-test splitting. In predicting the amount of reduction of the rent, the two models turned out to have a similar performance: their absolute mean errors are 0.1085 and 0.1119 for the random forest and the neural network, respectively. See Figure 2.

The neural network is fully connected with 21 nodes of input, followed by 3 hidden ReLu layers[9] of 256, 128, and 64 nodes, respectively, and one sigmoidal node for output, for a total amount of 46,849 parameters. The loss is the mean absolute error, trained by the Adam optimizer for 200 epochs and a minibatch size of 32. All other hyperparameters are default Keras values.

The random forest is a `sklearn.ensemble.RandomForestRegressor` with 100 estimators, a minimum samples split of 10, and cross-validation score with scoring = `neg_mean_absolute_error`[10].

The performances of the random forest, the linear model and the neural network are similar. For the explainability part described in the next paragraph, we used only the random forest. With respect to a neural network, a random forest is more easily optimized, more stable, and less opaque. Moreover, all the standard advantages that a neural network could have against a random forest disappear when using so few data points for training. This could possibly change when the feature extraction with NLP will be implemented, see Section 'Limitations and future developments'.

Explainability. We release all the code, models, and weights as open source. However, this does not mean that one can understand why a certain decision has been proposed to the judge simply by reading the code or looking at weights. This is because the models we used – random forests and neural networks – are intrinsically black box. Nevertheless, if we want to create a predictive justice framework that can be applied to real-life situations, then not being able to understand a decision is simply not acceptable. So, to overcome this limitation we used SHAP [Lundberg and Lee 2017], a state-of-the-art explainability technique that produces explanations, that is, arguments

[8] The names of the judges who consented to the mention are: Giovanna Bilò (Judge at the Court of Ancona, currently placed out of office at the Italian division of the European Court of Human Rights - Strasbourg), Anna Francesca Capone (Judge at the Court of Lecce), Luigi D'Orazio (Judge at the Court of Cassation), Giorgio Di Benedetto (President Emeritus at the Court of Sulmona), Francesca Di Donato (MOT at the Court of Napoli), Alessandro Di Tano (Judge at the Court of Ancona), Salvatore Grillo (Judge at the Court of Appeal of Bari), Rachele Dumella De Rosa (MOT at the Court of Napoli), Mario Fucito (Judge at the Court of Napoli), Martina Fusco (MOT at the Court of Napoli), D.ssa Emanuela Gallo (MOT at the Court of Napoli), Antonio La Catena (Judge at the Court of Foggia), D.ssa Annagrazia Lenti (Judge at the Court of Taranto), Claudia Malafronte (MOT at the Court of Napoli), Dr. Guglielmo Manera (Judge at the Court of Napoli), Aldo Marcaccio (MOT at the Court of Napoli), Chiara Maria Marcaccio (MOT at the Court of Napoli), Carlo Picuno (Judge at the Court of Audit), Eleonora Maria Velia Porcelli (Judge at the Court of Milano), Pasquale Angelo Spina (Judge at the Court of Castrovillari), Vincenzo Trinchillo (Judge at the Court of Napoli), Raffaele Viglione (Judge at the Court of Taranto).

[9] Readers not familiar with ReLU and activation functions in general can refer to Wiczew [2023] for a quick guide to activation functions.

[10] All the other default values can be seen at Scikit-Learn [2024].

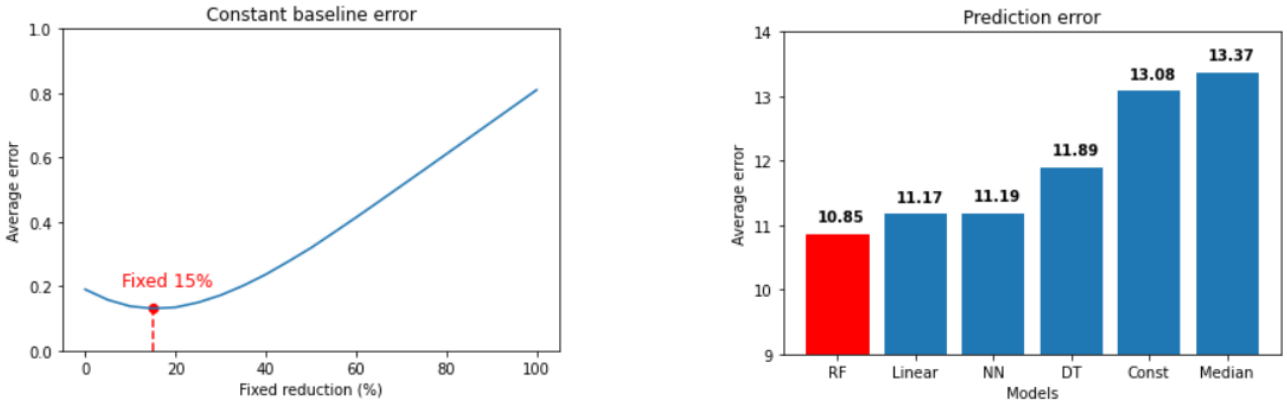

Figure 2: *Left: mean absolute errors of constant models. The best performance is achieved with a constant reduction of* 15%. *Right: histogram of mean absolute errors. From left to right: a random forest (RF), a linear model (Linear), a neural network (NN), a decision tree (DT), the best constant model (Const), and a model predicting the median value (Median). In this use case, with synthetic defense deeds and few data points, the performances of RF, Linear, and NN are similar.*

helping human users to understand the proposed decision. In particular, we measure the relative importance of each of the 25 features. In this way, we identify the factual elements that most affected the response provided by the AI. Not surprisingly, the features that turned out to be more relevant to the model decision are 'Percentage of loss of income to the tenant's total income', 'amount of support measures' and 'monthly rent amount'. Those alone contribute more than 71% to the final prediction, see Figures 6 in annex. Several other SHAP-derived explanations can be found in annex. These explanations can be used in real life to overcome objections regarding the opacity of AI.

We remark that we could have used a simple and interpretable by design model (a linear model or a decision tree), as shown by the baselines in Figure 2. However, we plan to use A.I.A.Co. to deal with much more complex use cases, that will very unlikely be described by a linear model[11].

Related work

How to use machine learning to estimate judges' decisions? This question is crucial in predictive justice, and can now be efficiently tackled thanks to the growing availability of legal big data. In the following, we list a few papers that in our opinion share some ground with our project. This list is by no means exhaustive.

In Chen and Eagel [2017], half a million asylum hearings rendered by 441 unique judges over a 32-year period, from 336 different hearing locations, were analyzed. The authors consider the binary task of granting / not granting asylum and use a random forest for classification. In Aletras et al. [2016], NLP is used to predict the binary outcome (violation yes/no) of cases tried by the European Court of Human Rights based solely on textual content. Masha Medvedeva et al. [2022] is a review of NLP methods applied to classification tasks in legal prediction. As in Chen and Eagel [2017], we use a random forest in the prototype, and we will use NLP as in Aletras et al. [2016] and Masha Medvedeva et al. [2022] in the next version of the software. However, we frame the problem as a regression task, instead of classification, because we need to predict a numerical quantity in $[0,1]$ (the reduction rate).

In the seminal paper by Kort [1957], the author develops a formula to extract numerical features that are then used to determine a decision boundary. This is not machine learning, but as far as we know it is one of the first contributions to the field of predictive justice and in some sense,

[11] A prototype of the software described in this paper is available at `https://vlab3.unich.it/aiaco`. All the source code, models, and training data will be available at `https://github.com/MistyDay86/A.I.A.Co`.

like us, deals with numerical quantities and not only with classification.

Then, Daniel Martin Katz [2012] coined the term 'quantitative legal prediction', trying to respond to questions like 'Do I have a case?, What is our likely exposure?, How much is this going to cost?, What will happen if we leave this particular provision out of this contract? and How can we best staff this particular legal matter?'. Our research goes in this direction by trying to answer the question: 'How much is the reduction of the rental fee?'. As a possible future development, when the NLP phase will be completed, we plan to extend the A.I.A.Co. framework to other quantitative legal questions, like, for example, 'How much is the maintenance allowance?'. See Section 'Limitations and future developments' for more details.

A recent interesting paper approaching quantitative legal prediction is Ilias Chalkidis et al. [2019b]. The authors use attention-based NLP models to predict the importance of a case on a scale from 1 to 4 in a regression task. To train and evaluate their models, the authors develop and release a very useful English legal judgment prediction dataset [Ilias Chalkidis et al. 2019a]. Similar to what we plan to do in the final version of our project, they use NLP; however, A.I.A.Co. is an end-to-end framework that will process legal documents starting from their original format (e.g. PDF), extracting both the input features and the label to be used during training. In this sense, when the NLP features extraction part will be completed, A.I.A.Co. will be an instance of self-supervised learning [LeCun and Misra 2021].

Validation survey

We designed a survey with the dual aim of evaluating both the impact on the justice system as well as the effectiveness of explanation methods as a tool to assist judges.

Survey structure

The survey is organized into five question panels, from **Q1** to **Q5** in the following.
Q1. The first panel is used to collect general personal data, their profession and degree of experience and request informed consent from participants.
Q2. In the second panel, the participants are presented with a synthetic defense deed and the corresponding decision predicted by the prototype model. Participants are asked if they are satisfied with the presentation of the case, with the quantity and quality of the information contained in the premises, and with their plausibility. As a final question they are asked to what extent they agree with the decision predicted by the model.
Q3. The third panel is devoted to the presentation of the three most informative features according to the SHAP model. Participants are asked to what extent they agree that the features presented are the most relevant to the case under discussion, what they think of the values assigned by SHAP, how much confidence they have in the predicted decision after having seen the SHAP explanation and, finally, how useful they find having the additional information about feature importance available in a real scenario.
Q4. In the fourth panel, we present three different counterfactual instances[12]. Each instance is derived from the original one by selecting the 3 most informative features presented in the previous panel as a set of actionable features, one at a time for each counterfactual explanation. For this, we use the BF counterfactual method in the *Fat Forensic Library*[13]. The BF method was chosen due to its ability to generate diverse counterfactuals, which is crucial for our validation survey. It allows us to present a range of possible scenarios to the participants, thereby enriching their understanding of the model's behavior. Participants are then asked to rate their confidence towards such counterfactual instances and, finally, how useful they find it to have a model available from which they can request counterfactuals such as those just described. See Figure 12 of

[12] In our A.I.A.Co. system, we employ the concept of counterfactuals as a means to provide explanations for the model's predictions. A counterfactual explanation describes a minimal change to the input features that would alter the model's prediction. In other words, it answers the question: 'What would have to be minimally different for the outcome to change?' This approach is particularly useful in understanding complex models, as it provides actionable insights for the users'.
[13] `https://fat-forensics.org/`.

the annex for an example.
Q5. In the final panel, we ask for an overall assessment of and the confidence in predictive justice models as seen in the case under discussion, leaving a blank field at the end of the survey to write comments, suggestions, and criticisms related to the use case.
In all quantitative questions, participants can rate on a scale from 1 to 5.

Hypothesis

The structure of the investigation reflects the following hypothesis we intend to address.

- **H1**. The synthetic defense deed is satisfactorily presented in its contents and in the order and structure of the information.
- **H2**. Participants generally agree with the predicted decision and, moreover, show greater confidence in that decision after having received the explanations.
- **H3**. Participants find counterfactual instances a useful aid tool for assisting in real scenarios.

Survey results

A total of 138 participants took part in the study. Participants signed up for the survey online after digitally signing a consent form followed by a short demographic survey. Aggregate demographic statistics of the participants are available in Figure 3. All participants have a legal background and, among them, as many as 67.3% belong to key categories for our survey, such as judges (16), legal scholars (46), and lawyers (31).

Figure 4a displays the participants' confidence in the model prediction both before and after viewing the explanations provided by SHAP. While the impact of receiving explanations is not appreciable across all categories, a significant increase in confidence is observed when the analysis is restricted to judges, the most significant category.

We performed a One-Way ANOVA test to statistically analyze these differences. The test yielded a p-value of 0.033 and an F-statistic of 2.70, indicating a statistically significant difference between the groups and thus rejecting the null hypothesis.

We also observe that the confidence level across all participant categories reaches a minimum rating of 3.0. This finding somewhat confirms hypothesis **H2**, particularly when considering the judges' category.

Figure 4b presents the participants' confidence in counterfactual instances. As with the previous analysis, judges show a markedly better reaction, while other categories still maintain a good rating level.

Applying the same One-Way ANOVA test to this data, we obtained a p-value of 0.042 and an F-statistic of 2.54, again indicating a statistically significant difference between the groups.

Furthermore, participants rated the usefulness of having counterfactuals available with an average score of 3.72 ± 0.59. This score is even more promising when restricted to judges (4.25 ± 0.41), while it reaches its lowest value of 3.50 ± 0.50 among law students. This analysis addresses hypothesis **H3**.

Hypothesis **H1** is addressed by examining the responses to the first part of question **Q2**. Participants gave an average score of 3.62 ± 0.51, with judges' scores rising to 4.37 ± 0.35.

All the analyses described above demonstrate that judges consistently achieve the highest scores across all questions posed to the participants.

Impact to justice

Discretionality

The framework developed in this project serves as a reference tool for Judges to estimate the reduction of rental fees in favor of tenants. The software/prototype described in the paper does not aim to predict all the *force major* situations, but rather intends to provide an equitable solution for renegotiation in the event of contingencies, regardless of their nature.

This mechanism has the potential to promote greater uniformity in judgments and enhance legal certainty. A similar system is already being used in Italy to determine compen-

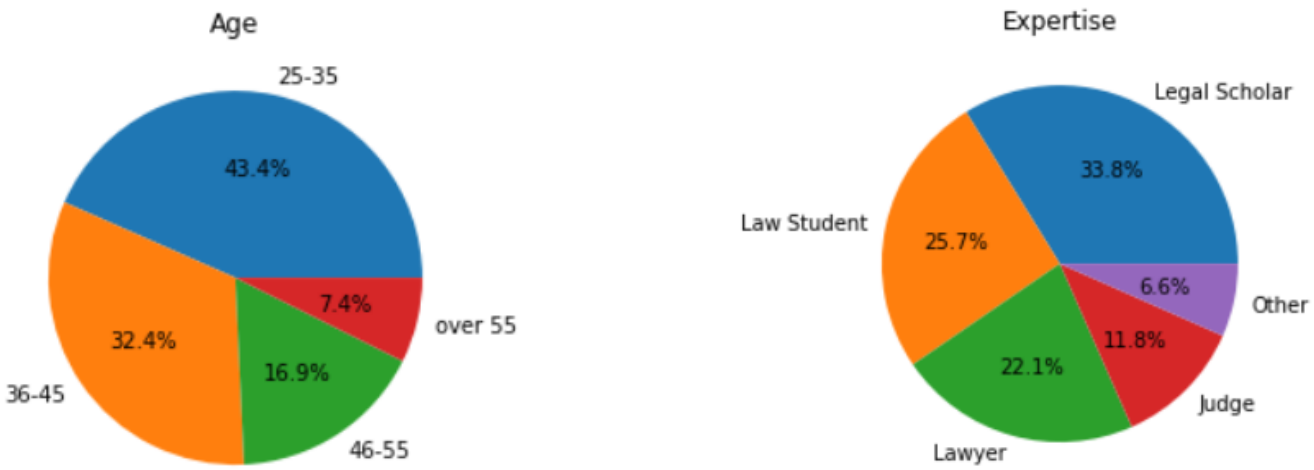

Figure 3: *Aggregate demographic statistics of the participants.*

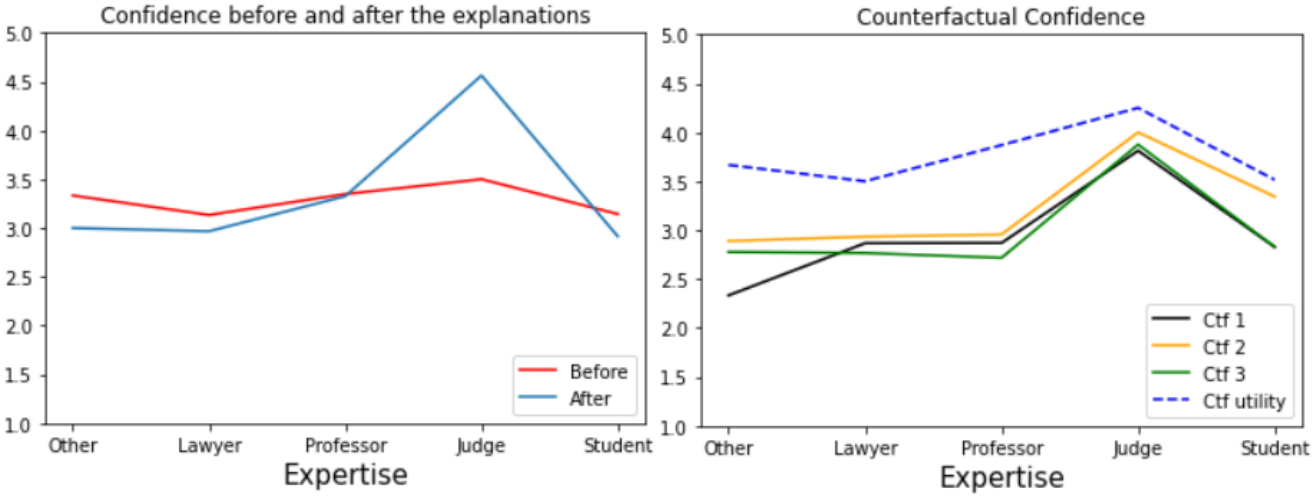

(a) **(b)**

Figure 4: *Participants' confidence in (a) the model prediction before and after looking at the SHAP explanation and in (b) the counterfactual instances.*

sation for non-pecuniary damages, as exemplified by the '*Tabelle di Milano*', developed by the Observatory of the Court of Milano and used throughout the Italian national territory [Osservatorio sulla giustizia civile di Milano 2021]. However, the aforementioned tables do not eliminate the discretion of the judge, who can always 'customize' the decision proposed by the algorithm while adequately justifying his decision.

Indeed, in our case the judge also has to take heed of possible biases or potential software limitations and, where appropriate, make corrections for those. Since the responsibility ultimately rests with the judge, this remains necessary.

Moreover, law serves as a lens through which reality is viewed, and judges are required to interpret the law in light of historical, sociological, and political changes at the time of decision-making. These aspects are not easily grasped by machines, as 'predictive justice' is derived from already published judicial decisions. Therefore, the presence of the Judge as a natural person becomes essential to reevaluate the decision proposed by the algorithm[14]. Therefore, the AI suggestion should only serve as a 'second opinion', that helps legitimize the introduction of the framework despite any potential weaknesses. This perspective aligns with the prevailing opinion (see Allhutter et al. [2020]) and is also lastly reflected in recital no. 61 of the A.I. Act[15], which states that 'the use of AI tools can support the decision-making power of judges or judicial independence, but should not replace it: the final decision-making must remain a human-driven activity'.

On the other hand, human decision-making is also subject to biases. As Holmes points out, legal certainty is never given because it depends on the foreseeability of human action and human interpretation (that is, in the civil proceedings, the judge), both of which are fundamentally uncertain. On this point, see also Alaire [2016], Hildebrandt [2018], D. M. Katz et al. [2017], and Kurzweil [2005]. Moreover, different judges may decide on the same issue in dissimilar ways, so an automated decision could replicate human bias.

It would be necessary/useful because, currently, Italian judges are tasked with determining rental fee reductions based on equity when they become excessively burdensome. One criticism of this approach (i.e. of an equity-based decision) is that judges typically decide without specific parameters to guide their decisions (our prototype offer a suggestion that takes into account similar rulings made by other judges).

For these reasons, a man-machine integrated system allows us to overcome both the human and the artificial limitations, aiming to propose a decision that is as fair as possible. More generally, on the importance of the man-machine integration see Pasquale [2016, 2020]. Consequently, it may also help avoid the application of rules on fully automated decision-making systems, as outlined in the EU Reg. 2016/679, commonly known as the G.D.P.R.

Currently, the use of A.I.A.Co. is not intended to be mandatory for judges. The discretion of the judge should not be eliminated, similar to the case of the aforementioned '*Tabelle di Milano*'. The judges could therefore deviate from the proposed solution if they do not share the application's assumptions. In fact, the framework would be applicable only if shared by the judge, and this is always in view of a judge-centric justice.

Accountability

Of course, even if the judge chooses to use the proposed solution, he still has the responsibility to justify the pro-

[14] Generally, on the issues raised by predictive justice in the Italian context, see Dalfino [2018], Gabellini [2019], Mattera [2018], Rulli [2018], and Viola [2018].

[15] See Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139, and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). Furthermore, recital no. 61 considers that AI systems intended for the administration of justice and democratic processes (such as those intended to be used by or on behalf of a judicial authority), as well as those intended to be used by alternative dispute resolution bodies (when the outcomes of alternative dispute resolution proceedings have legal effects on the parties), should be classified as high-risk, 'considering their potentially significant impact on democracy, the rule of law, individual freedoms as well as the right to an effective remedy and to a fair trial'.

vision and the solution adopted in the specific case. The availability of the source code and models and the used explainability mechanism do not replace the judge's duty to provide justification. Instead, it enables the parties involved and the judge himself to understand how and why a particular decision was proposed. In particular, the multidisciplinary nature underlying the design of A.I.A.Co. – a design that requires not only legal skills but also technical, mathematical, and statistical ones – does not exempt the judge from accompanying the technical description of the framework with a clear explanation for the decision, making it readable and comprehensible to the parties. This allows the parties to understand the underlying assumptions and review the results. For more on the significance of collaboration between legal professionals and computer scientists in developing technology that protects human rights, see also Hildebrandt [2016].

Transparency and clarity

The mechanism through which the decision is made must be known to the parties in advance, in alignment with the principles of transparency and clarity. This includes awareness of the authors of the decision, the procedure used for its development, the decision-making mechanism (including the assigned priorities in the evaluation process), and the selected relevant data. On this point, see also Diver [2020] and Goldoni [2015]. Parties should be able to verify that the criteria, conditions, and results of the procedure adhere to the requirements and purposes established by the law. This also allows parties to ascertain if the applied methods and rules are clear. Furthermore, downstream verification of the results and their logical correctness must be ensured (on this point, see also Cons. State n. 8472/2019).

Every component of A.I.A.Co. is open source, precisely to allow the parties and the judge to perform a thorough analysis before deciding whether to use it. This approach helps address concerns about the inherent opacity of the used models. For more on the issue of AI system opacity, refer to Hildebrandt [2018] and Pasquale [2005].

Duration of proceedings

Furthermore, the judicial use of A.I.A.Co. in this type of trial may have an impact on the duration of proceedings, which is a significant issue, especially in Italy. It can expedite decision-making by assisting the judge in identifying the criteria for parameterizing the decision and providing a tool that proposes the decision itself. This aligns with the principle of a 'fair trial' as granted by art. 111 of the Italian Constitution and by art. 6 ECHR. The European Court of Human Rights repeatedly specified that, among the guarantees of a fair trial, there is also respect for the reasonable length of the trial, as an instrument to guarantee the efficiency and credibility of justice (see European Court of Human Rights, 24 October 1989, H v. France). On the other hand, during the non-litigation phase, A.I.A.Co. functions as a tool to reduce the number of cases brought before the Courts, with inevitable economic benefits for the parties involved, because parties can anticipate and share the A.I.A.Co. decision before seeking the involvement of a professional judge. In such cases, it would be desirable for the algorithm to be utilized by specialized entities, such as Arbitration Chambers, mediation or conciliation bodies and other organizations operating within various Councils of Bar Associations, Chambers of Commerce, etc. In this manner, this predictive framework has the potential to facilitate access to justice from a democratic justice perspective and to ensure the effective protection of human rights. This aligns with Goal 16 'Peace, justice, and strong institutions' of the 2030 Agenda for Sustainable Development, adopted in September 2015 by the governments of UN member countries.

Appeals system

Regarding the appeals system provided by the Legislator, it should be noted that the algorithm does not alter the existing regulations. For judicial measures that base their decision on A.I.A.Co. predictions, the ordinary appeal system would remain applicable. However, if A.I.A.Co. is used as an out-of-court remedy and the parties do not agree on the proposed solution, an agreement will not be reached, and the parties will have to seek the involvement of a professional judge, similar to the current systems of Alternative Dispute Resolution.

Limitations and future developments

One of the future research directions arising from this Project is the promotion of Artificial Legal Intelligence in service of the rule of law. In this way, as already stated by Hildebrandt, we could pass from 'legal by design' to 'legal protection by design' [Hildebrandt 2018].

All the participants involved in the Project have already acknowledged the potential of the device. Specifically, almost all contacted Magistrates have expressed interest in experimenting with and utilizing the final prototype in their own Judicial Offices.

It should be noted that the current limitation of the A.I.A.Co. project lies in the absence of the initial NLP feature extraction component. Nevertheless, this prototype provides crucial insights for the future development of the NLP model, such as SHAP suggesting that the salient features to be extracted are likely to be fewer than 10.

One notable limitation of the current A.I.A.Co. proof-of-concept implementation is the distribution shift between the synthetically generated cases and real-world cases. To mitigate this issue as much as possible, we have collaborated with domain experts in selecting the training data, as detailed in the 'Generation of training data' paragraph in Subsection 'Methods and results'. Furthermore, we have conducted an additional verification through question **Q2** in the survey to confirm the effectiveness of this mitigation strategy. In the final version of A.I.A.Co., the training data will be derived from the actual data distribution rather than relying on synthetic data, we will elaborate on this in the next paragraph. This approach should significantly reduce the impact of the distribution shift, making the system more robust and accurate in handling real-world cases.

Currently, A.I.A.Co. lacks the capability to self-evaluate input quality and output uncertainty. These issues could be mitigated by several strategies. To self-evaluate input quality, one could use for instance anomaly detection techniques, where outliers in new inputs are identified based on 'normal' data patterns learned from historical inputs. Data consistency checks could ensure that new data aligns with historical patterns, flagging significant deviations as potential quality issues. A separate model could also be trained to recognize out-of-distribution data, assigning each new input a quality score or accept/reject label. To self-evaluate the outputs uncertainty, one approach could be to use probabilistic models, which can provide a measure of uncertainty along with their predictions. Alternatively, techniques such as bootstrapping or ensemble methods could be used to estimate the variability in the predictions.

In the future we will implement feature extraction using NLP techniques leveraging the data already available at the Italian Ministry of Justice through the so-called PCT (Telematic Civil Process). In particular, we will use the Italian judges' statements in authentic civil judgments, already published in the so-called Merit Database[16], which is constantly updated (as new judgments are added[17]). Outside the Italian national panorama, quantitative analysis systems for jurisprudence are already in use (e.g. U.S.A., Belgium, Czech Republic, France, Germany, and Israel), so this system will be easily exportable abroad. On this point, see [M. Medvedeva et al. 2020]. The ISTI-CNR of Pisa ([18]) and the PI School ([19]) have expressed interest in collaborating on the future NLP step.

Although this predictive framework was initially designed to address specific issues that stemmed from the Covid-19 lockdown, the acquired knowledge, the produced model and the research outcomes can be directly transfered to other civil issues. The predictive framework could, for example, also be applied to the following domains: family law (e.g. to determine the amount of the maintenance or di-

[16] That is free and freely accessible at `https://bdp.giustizia.it/login`.

[17] This data will expand also thanks to the Italian National Recovery and Resilience Plan, the so-called N.R.R.P.: as part of the 'Digitization, Innovation, and Security in the P.A.' mission funded by the N.R.R.P., the Ministry of Justice aims to digitize approximately 10 million hybrid and paper Court case files for proceedings spanning 2016-2026. More information are available at `https://padigitale2026.gov.it/misure/`.

[18] `https://isti.cnr.it/en/`.

[19] `https://picampus-school.com/`.

vorce obligation); the calculation of the compensation for damages, including in the insurance field; executive and insolvency procedures (e.g. as regards to the determination of allocation plans; moreover the N.R.R.P. reforms the negotiated composition for the solution of business crisis, establishing a software capable of verifying the sustainability of debts and automatically drawing up a recovery plan for liabilities below a certain threshold, see art. 30 d.l. n. 118/2021, conv. l. n. 147/2021, and art. 30-ter and 30-quinquies of d.l. n. 152/2021, conv. l. n. 233/2021); renegotiation of business contracts as part of the proceedings of the negotiated composition of the business crisis pursuant to art. 10 of d.l. n. 118/2021, conv. l. n. 147/2021.

Similar applications are already being used: in Perù a Judge recently used the known software 'ChatGPT' to determine the maintenance allowance for a little girl[20] and in Colombia the Court made use of the same program to resolve a lawsuit initiated by a mother seeking exemption from payment for hospital appointments and treatment[21].

Furthermore, the potential of A.I.A.Co. expands when considering international projects aiming to translate rules into self-executing codes. Numerous sources have explored this area (*ex multis*, Diver [2020], Merigoux et al. [2021], and Mohun and Roberts [2020]). With the adoption of such projects, the utilization of A.I.A.Co. becomes even more feasible.

Finally, an interesting question has been raised by a referee. Since we are dealing with situations that are by nature exceptional, this exceptionality seems antithetical to predictive automated calculations. As mentioned above, the software/prototype described in the paper does not aim to predict *force majeure*/exceptional situations: it intends only to offer an equitable solution for renegotiation in case of contingencies. We do not exclude the possibility of exploring the contradiction between exceptional situations and predictive automated calculation in future research.

Lastly, we remark that due to the missing NLP feature extraction part, this paper should be considered a proof-of-concept, rather than the presentation of a robust deployable system. This proof-of-concept supports the possibility of using Artificial Intelligence to resolve certain legal issues, as was extensively described in this section.

Conclusions

This predictive framework has significant potential and differs from other similar Italian initiatives[22]. In our project AI system is put to use, applied or consulted in concrete-decision making, while ensuring that the final assessment remains with human judges. Therefore, our project can represent an important advancement in legal prediction research and can potentially be used, with minor adaptations, to a wide range of situations, as mentioned above.

Despite the broad potential of A.I.A.Co., the decision – as mentioned in Section 'Impact to justice' – must remain judge-centric, that is, subject to human control, as required by the anthropocentric approach. The entire research project was conceptualized from this perspective, which was also at the core of the European strategy on AI, as expressed in the White Paper on AI of 19/02/2020 (2020/65/COM), and of the so-called A.I. Act (Regulation (EU) no. 2024/1689) of 13 June 2024[23]. The irreplaceable nature of the judge-natural person in the use of AI systems was never called into question. This approach is also in accordance with the provisions of the Italian National Research Program-P.N.R. 2021-2027, which focuses on human-centered innovation.

The NLP version of A.I.A.Co. could assist not only the judges but also their auxiliaries (such as, for example, the professional delegate in executive procedures), the

[20] Juzgado Transitorio Dd Familia De San Juan De Miraflores (Perù), Judge Frank Flores García, ruling in Exp. No. 00052-2022-18-3002-JP-FC-01 of 27/03/2023.

[21] Juzgado Primero Laboral di Cartagena de Indias (Colombia), ruling no. 032 of 30/01/2023.

[22] See in particular the researches from the University of Brescia, the University Sant'Anna of Pisa and the *Alma Mater Studiorum* of Bologna, available at `https://giustiziapredittiva.unibs.it/home`, `https://site.unibo.it/cross-justice/en/project-results/tools` and `https://www.predictivejurisprudence.eu/`, respectively.

[23] On this point see also the work of the European Commission for the Efficiency of Justice - CEPEJ available at `https://www.coe.int/en/web/cepej`.

parties involved and the other actors of the out-of-court phase (lawyers, mediators, conciliators, experts, insurance companies, *etc.*), aiming for an overall reduction of disputes. For more on the use of technology to support lawyers in general, see I. Chalkidis et al. [2019] and Steer [2021].

A.I.A.Co. exhibits greater flexibility than other similar AI tools, overcoming the known challenges in the field of predictive justice, without questioning the irreplaceable nature of the judge-natural person in the use of AI systems. Moreover, A.I.A.Co. is aligned with the objectives of modernizing and streamlining the judicial system pursued by the recent reform of the Italian civil process (l. n. 206/2021, known as 'Cartabia reform'), by the establishment of the '*Ufficio per il processo*', which primarily aims to ensure reasonable duration of proceedings through the use of technology and innovation processes in judicial offices[24]. Additionally, the Italian National Recovery and Resilience Plan seeks to facilitate dispute resolution and promote the use of technology to improve process efficiency and reduce the workload of Magistrates. The obtained results can serve as a valuable and immediate tool for Judicial Offices as well as other Alternative Dispute Resolution methods (e.g. the Arbitration Chambers or the mediation or conciliation bodies and others operating within the various Councils of the Bar Associations, Chambers of Commerce, *etc.*).

Ultimately, this project demonstrates that man-machine integration cannot be ignored in order to achieve fair justice. AI utilization in predictive justice systems should neither be demonized nor absolutized, avoiding the dangerous extremes of techno-optimism or techno-pessimism. In this vein, the prior interconnection between the legal and technological realms is necessary, which must dialogue and complement each other.

References

B. Alaire. 2016. 'The Path of the Law: Towards Legal Singularity'. *University of Toronto Law Journal*, 66, 4.

Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preotiuc-Pietro, and Vasileios Lampos. 2016. 'Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective'. *PeerJ Comput. Sci.*, 2, e93.

D. Allhutter, F. Cech, F. Fischer, G. Grill, and A. Mager. 2020. 'Algorithmic Profiling of Job Seekers in Austria: How Austerity Politics Are Made Effective'. *Frontiers in Big Data*, 3.

A. M. Benedetti. 2020. *Stato di emergenza, immunità del debitore e sospensione del contratto*. `www.giustiziacivile.com`. (2020).

G. Carapezza Figlia. 2020. 'Coronavirus e locazioni commerciali. Un diritto eccezionale per lo stato di emergenza?' *Actualidad Juridica Iberoamericana*.

I. Chalkidis, I. Androutsopoulos, and N. Aletras. 2019. 'Neural Legal Judgment Prediction in English'. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4317–4323.

Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019a. *Dataset for paper 'Neural Legal Judgment Prediction in English'*. https://archive.org/download/ECHR-ACL2019. (2019).

Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. July 2019b. 'Neural Legal Judgment Prediction in English'. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Florence, Italy, (July 2019), 4317–4323.

Daniel L. Chen and Jess Eagel. 2017. 'Can Machine Learning Help Predict the Outcome of Asylum Adjudications?' In: *Proceedings of the 16th Edition of the International Conference on Articial Intelligence and Law* (ICAIL '17). Association for Computing Machinery, London, United Kingdom, 237–240.

Giuseppe Conte. 2019. *Disegni di legge atto Senato n. 1151*. `https://www.senato.it/leg/18/BGT/Schede/Ddliter/51488.htm`.

D. Dalfino. 2018. 'Creatività e creazionismo, prevedibilità e predittività'. *Il Foro Italiano*, 12, 385–393.

L. Diver. July 2020. 'Digisprudence: the design of legitimate code'. *LawArXiv*, (July 2020). `osf.io/preprints/lawarxiv/nechu`.

A. A. Dolmetta. 2020a. *Il problema della rinegoziazione (ai tempi del coronavirus)*. `www.giustiziacivile.com`. (2020).

A. A. Dolmetta. 2020b. *Locazione di esercizio commerciale (o di studi professionali) e riduzione del canone per "misure di contenimento" pandemico*. `blog.ilcaso.it/libreriaFile/1210.pdf`. (2020).

European Law Institute. 2020. *ELI Principles for the COVID-19 Crisis*. `https://www.europeanlawinstitute.eu/fileadmin/user_upload/p_eli/Publications/ELI_Principles_for_the_COVID-19_Crisis.pdf`.

A. Federico. 2020. 'Misure di contenimento della pandemia e rapporti contrattuali'. *Actualidad Juridica Iberoamericana*.

E. Gabellini. 2019. 'La "comodità nel giudicare": la decisione robotica'. *Rivista trimestrale di diritto e procedura civile*.

A. Gentili. 2003. 'La replica della stipula: riproduzione, rinnovazione, rinegoziazione del contratto'. *Contr. e impr.*

M. Goldoni. 2015. *The Politics of Code as Law: Toward Input Reasons*. Information and Law in Transition: Freedom of Speech, the Internet, Privacy and Democracy in the 21st Century. Liber.

M. Grondona. 2020. 'Dall'emergenza sanitaria all'emergenza economica: l'eccessiva onerosità sopravvenuta tra buona fede e obbligo di rinegoziazione'. *Actualidad Juridica Iberoamericana*.

[24] See art. 16-*octies*, d.l. n. 179/2012, conv. l. n. 221/2012 and `https://www.giustizia.it/giustizia/it/mg_2_9_2.page`.

M. Hildebrandt. 2018. 'Law as computation in the era of artificial legal intelligence: Speaking law to the power of statistics'. *University of Toronto Law Journal*, 68, 1, 12–35. eprint: https://doi.org/10.3138/utlj.2017-0044.

M. Hildebrandt. 2016. 'Law as Information in the Era of Data-Driven Agency'. *The Modern Law Review*, 79, 1, 1–30. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-2230.12165.

International Chamber of Commerce. 2020. *ICC Force Majeure Clause*. https://iccwbo.org/news-publications/icc-rules-guidelines/icc-force-majeure-and-hardship-clauses/.

D. M. Katz, M. J. Bommarito II, and J. Blackman. Apr. 2017. 'A general approach for predicting the behavior of the Supreme Court of the United States'. *PLOS ONE*, 12, 4, (Apr. 2017), 1–18.

Daniel Martin Katz. 2012. 'Quantitative Legal Prediction – or – How I Learned to Stop Worrying and Start Preparing for the Data Driven Future of the Legal Services Industry'. *Emory law journal*, 62, 909.

Fred Kort. 1957. 'Predicting Supreme Court Decisions Mathematically: A Quantitative Analysis of the "Right to Counsel" Cases'. *American Political Science Review*, 51, 1–12.

Ray Kurzweil. 2005. *The Singularity is near: When Humans Transcend Biology*. Viking, New York.

Yann LeCun and Ishan Misra. 2021. *MetaAI blog post'*. https://ai.facebook.com/blog/self-supervised-learning-the-dark-matter-of-intelligence. (2021).

N. Lipari. 2003. 'Valori costituzionali e procedimento interpretativo'. *Rivista trimestrale di diritto e procedura civile*.

S. M. Lundberg and Su-In Lee. 2017. 'A unified approach to interpreting model predictions'. In: *Advances in Neural Information Processing Systems 2017*, 4768–4777.

F. Macario. 1996. *Adeguamento e rinegoziazione dei contratti a lungo termine*. Jovene, Napoli.

F. Macario. 2020a. *Per un diritto dei contratti più solidale in epoca di "coronavirus"*. www.giustiziacivile.com. (2020).

F. Macario. 2020b. 'Sopravvenienze e rimedi al tempo del "Coronavirus": interesse individuale e solidarietà'. *Contratti*.

R. Mattera. 2018. 'Decisione negoziale e giudiziale: quale spazio per la robotica?' *La nuova giur. civ. comm.*, II.

M. Medvedeva, M. Vols, and M. Wieling. June 2020. 'Using Machine Learning to Predict Decisions of the European Court of Human Rights'. *Artif. Intell. Law*, 28, 2, (June 2020), 237–266.

Masha Medvedeva, Martijn Wieling, and Michel Vols. 2022. 'Rethinking the field of automatic prediction of court decisions'. *Artificial Intelligence and Law*.

D. Merigoux, N. Chataing, and J. Protzenko. 2021. 'Catala: A Programming Language for the Law'. *Proc. ACM Program. Lang.*, 5, ICFP, 1–29.

J. Mohun and A. Roberts. Oct. 2020. *Cracking the code*. OECD Working Papers on Public Governance 42. (Oct. 2020).

Osservatorio sulla giustizia civile di Milano. 2021. *Tabelle milanesi per la liquidazione del danno non patrimoniale*. https://tribunale-milano.giustizia.it/cmsresources/cms/documents/Tabelle_milanesi_Danno_non_patrimoniale_ed._2021_2.pdf.

F. Pasquale. 2016. 'Automating the Professions: Utopian Pipe Dream or Dystopian Nightmare?' *Los Angeles Review of Books*. lareviewofbooks.org/article/automating-the-professions-utopian-pipe-dream-or-dystopian-nightmare.

F. Pasquale. 2020. *New Laws of Robotics: Defending Human Expertise in the Age of AI*. Belknap Press.

F. Pasquale. 2005. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press.

P. Perlingieri. 2017. *"Controllo" e "conformazione" degli atti di autonomia negoziale*. Rassegna di diritto civile. Vol. 1 (2017). Edizioni Scientifiche Italiane.

P. Perlingieri. 2001. *Equilibrio normativo e principio di proporzionalità nei contratti*. Rassegna di diritto civile. Vol. 3. Edizioni Scientifiche Italiane.

P. Perlingieri. 2011. *Il "giusto rimedio" nel diritto civile*. Il giusto processo civile. Vol. 6 (2011). Edizioni Scientifiche Italiane.

P. Perlingieri. 2020a. *Il diritto civile nella legalità costituzionale secondo il sistema italo-europeo delle fonti*. Vol. IV. Edizioni Scientifiche Italiane.

P. Perlingieri. 2020b. 'Principio personalista, dignità umana e rapporti civili'. *Lezioni*, III.

A. Pisu. 2021. 'Pandemia, locazioni commerciali e giustizia contrattuale: l'incerta avanzata dell'obbligo legale di rinegoziazione'. *Resp. civ. e prev.*, 4, 1295.

I. Prisco. 2021. *Sull'allocazione convenzionale del rischio contrattuale*. Rassegna di diritto civile. Vol. 3 (2021). Edizioni Scientifiche Italiane.

E. Rulli. 2018. 'Giustizia predittiva, intelligenza artificiale e modelli probabilistici. Chi ha paura degli algoritmi?' *Analisi Giuridica dell'Economia*, 2, 533–546.

Scikit-Learn. 2024. *API of RandomForestRegressor*. https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestRegressor.html.

G. Sicchiero. 2002. 'La rinegoziazione'. *Contr. e impr.*

P. Sirena. 2020. 'Eccessiva onerosità sopravvenuta e rinegoziazione del contratto: verso una riforma del codice civile?' *Jus*, 205.

S. Steer. 2021. 'Westlaw Edge UK: Three Traits that Make for an Unrivalled Lawyer'. *Legal Information Management*, 21, 2, 79–87.

L. Viola. 2018. *Interpretazione della legge con modelli matematici. Processo, a.d.r., giustizia predittiva*. Diritto Avanzato.

Daniel Wiczew. 2023. *Demystifying Activation Functions in Neural Networks: A Guide for Beginners*. https://www.linkedin.com/pulse/demystifying-activation-functions-neural-networks-guide-daniel-wiczew-dceae/.

A reply: Sentencing guidelines 2.0?

Liane Huttner • Université Paris-Saclay, huttner.liane@gmail.com

This paper introduces an AI-based predictive framework designed to assist judges in the determination of commercial rents during theCovid-19 pandemic. A.I.A.Co was created by an interdisciplinary team of lawyers and computer scientists, it is open source and outputs are explainable. As such, A.I.A.Co is a good example of transparent and explainable AI, in line with recent trends in the field. These approaches should be encouraged, and I strongly share the authors' opinions on the necessity for transparency and accountability in AI systems.

There are, however, two main legal questions that have not been fully addressed in the paper and that require further consideration. The first relates to the use of A.I.A.Co as an advisory tool, that should not replace the judge. This is key for the authors, as they point out multiple times: 'the ultimate responsibility for the decision rests with the judge', 'the AI suggestion should only serve as a "second opinion"', 'the use of A.I.A.Co is not intended to be mandatory for judges'. I wholeheartedly agree with this goal. It is an essential point in almost all digital laws, from the GDPR to the AI Act. Important decisions must be taken by humans, for humans. But stating that a tool should only be used as an advisory opinion is far from ensuring that it actually is used in such a way. The influence of AI systems (and even of algorithms or data processing that do not qualify as AI) on human decisions is a difficult and controversial aspect of digital law. It is generally agreed that there are two categories of automated decision-making tools: those that allow fully automated decisions, and those that serve as a partial support for human decision-making. This is the difference between fully automated decision-making and partially automated decision-making. This distinction does not refer to the complete absence of human judgement, but rather to its importance in the decision-making process. The distinction is far from straightforward. It is difficult to understand with certainty how any tool is used in a decision-making process[25]. We can't read a judge's mind – or anyone's, for that matter, if my sources are correct. This is why scholars and institutions have proposed various ways of assessing the use of technological artefacts as a basis for decision-making [Information Commissioner's Office 2018, p. 9; Brennan-Marquez et al. 2019, p. 750; Binns and Veale 2021, p. 320]. In my opinion, however, it is only through the design of a system that one can ensure it is exclusively used as an advisory opinion. This design must guarantee that it is impossible for the system to be used in a different way[26]. In other words, for a tool to be used in an exclusively advisory capacity, the need for human judgement must be embedded in the algorithm itself.

This is particularly true in cases such as A.I.A.Co. The authors correctly refer to the issue of the duration of proceedings. History has shown that when machines set the pace, humans must adapt. The opposite occurs rarely. If A.I.A.Co functioned properly and was indeed used in courts, I think it is highly unlikely that judges would take the time to review each case as they might have done without the tool. This is probable even when explanations are available [Yacoby et al. 2022]. Though not necessarily a bad thing in itself, this is nonetheless at odds with the asserted goal of A.I.A.Co. In other words, the authors do not specify what limits A.I.A.Co as a tool that can only be used as a partial basis for the decision.

This is also why regulations such as article 22 GDPR have been adopted. It is a shame that the authors seek to 'avoid the application of rules on fully automated decision-making systems,' since these rules have been created ex-

[25] See for example the conclusions of Advocate General Pikamäe in the Schufa case: CJEU, C-634/21, OQ v Land Hesse, 16/03/2023, § 42 and 43.

[26] A similar approach exists in the EU AI Act proposal: Article 14 requires that human oversight is identified and built into the high-risk AI system by the provider before it is placed on the market or put into service.

actly for such cases. If A.I.A.Co were to be used by judges, measures such as a right to obtain human intervention and a right to contest the decision should be the net minimum to protect litigants. Would it not be better to simply acknowledge that A.I.A.Co could be used as the sole basis of the decision, and therefore ensure that safeguards are in place?

The second legal issue that I would like to discuss is the mere existence of a tool like A.I.A.Co. Sentencing guidelines and models for compensation calculations have existed for decades. They are not always complex systems that use AI. They can be straightforward decision trees, or a few lines in a law that restrict the judge's power of appreciation. It should be said, though, that even in their simplest form, such guidelines are controversial.

In the abundant literature about judicial guidelines, the same types of arguments are often made [Ashworth and Roberts 2013]. Proponents believe that guidelines lead to uniformity, legal certainty, and consistency. Critics point out the threat to judicial independence, individualization of sentencing, and the more nebulous 'art of judging' [Gerry-Vernières 2019]. Finding a balance is delicate. The power of appreciation of judges, and their freedom, is at the core of the rule of law. It guarantees that all litigants can have their cases examined by a judge. This is the fundamental right of access to courts. Judicial guidelines, even when they are not mandatory, interfere with this right.

Furthermore, even when such guidelines are adopted by legislators, they are always at risk of being deemed illegal by courts. This happened recently in Italy. The European Committee of social rights decided that the predefined compensation mechanism for victims of unlawful dismissals violated article 24 of the European social charter because it did not 'make it possible in all cases of dismissal without valid reason to obtain appropriate redress proportionate to the damage suffered and apt to discourage employers from resorting to unlawful dismissal'[27].

This question is even more important in light of the envisioned direction of A.I.A.Co. The authors wish to use the model and the research outcome in other areas of law, such as family law, damages, insolvency, and business contract law. Each topic will pose specific problems, but potential issues are particularly noticeable in family law and damages. Guidelines for alimony have been discussed for decades, and it is still not certain whether they are useful or not [Willick 2014]. Calculation for damages must obey the principle of full compensation, and this sometimes requires judges to put a price on loss and anxiety. The A.I.A.Co research project does not fully acknowledge how delicate this is.

A.I.A.Co admirably takes into account many of the pressing issues in AI and law. If it bettered addressed the two legal concerns that I have outlined here, it could prove a useful tool for decision-making.

References

Andrew Ashworth and Julian V Roberts, eds. . 2013. *Sentencing guidelines: Exploring the English model*. OUP Oxford.

Reuben Binns and Michael Veale. 2021. 'Is that your final decision? Multi-stage profiling, selective effects, and Article 22 of the GDPR'. *International Data Privacy Law*, 11, 4, 319–332.

Kiel Brennan-Marquez, Karen Levy, and Daniel Susser. 2019. 'Strange Loops'. *Berkeley Technology Law Journal*, 34, 3, 745–772.

Stéphane Gerry-Vernières. Nov. 2019. *La barémisation de la justice*. Research Report 216.11.04.23. GIP Mission de recherche Droit et Justice, (Nov. 2019).

Information Commissionner's Office. 2018. *Detailed Guidance on Automated decision-making and profiling*. `https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/individual-rights/automated-decision-making-and-profiling/`.

Marshal S Willick. 2014. 'A Universal Approach to Alimony: How Alimony Should Be Calculated and Why'. *J. Am. Acad. Matrimonial Law.*, 27, 153.

Yaniv Yacoby, Ben Green, Christopher L Griffin Jr, and Finale Doshi-Velez. 2022. '"If it didn't happen, why would I change my decision?": How Judges Respond to Counterfactual Explanations for the Public Safety Assessment'. In: *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*. Vol. 10, 219–230.

[27] ECSR, Confederazione Generale Italiana del Lavoro (CGIL) v. Italy, Complaint No. 158/2017, 11/09/2019.

Author's reponse

Maurizio Parton and others

The reviewer points out that it would be difficult to ensure that the tool is actually – and only – used to get a second opinion. The review concludes by suggesting that it might better to acknowledge that there is a genuine risk that A.I.A.Co. would be used as the sole basis for decision-making and that it would actually make more sense to anticipate to this by providing safeguards. In reply to this suggestion, we would like to re-emphasize our belief that 'Important decisions must be made by humans, for humans' (as reiterated by the same review). This sentiment is in line with the recent European legislation, referred to in the article, and precisely with the recital no. 61 of the A.I. Act[28], which states that 'The use of AI tools can support the decision-making power of judges or judicial independence, but should not replace it: the final decision-making must remain a human-driven activity'. An additional risk to A.I.A.Co. being used as the sole basis for decision-making is the fact that judges might also be unduly influenced by the algorithm's decision. To address this concern, the proposed model still requires the judge's motivation for the judicial ruling and preserves their discretion. Judges can always deviate from the algorithm's solution by providing adequate justification. Moreover, it is crucial to recognize that human decision-making is also subject to bias. In addition, different judges may reach different conclusions when faced with the same issue and judges may be influenced by each other's decisions. Furthermore, the review states that the article highlights the existence of sentencing guidelines and compensation calculation models for decades. For the review 'Calculation for damages must obey the principle of full compensation and this sometimes requires judges to put a price on loss and anxiety'. Despite criticism that they are inflexible and do not take into account the nuances of individual cases, potentially leading to unjust or inappropriate decisions, the use of guidelines to promote consistency and predictability in judicial decisions is widespread in various legal domains. In addition to the mentioned '*Tabelle di Milano*', which are used for non-pecuniary damage compensation, we refer to the tables provided by the Italian Private Insurance Code for non-pecuniary damages resulting from traffic accidents, or in the field of family law, the 'Family Protocols' established by Bar Associations and related Courts, which provide guidance on economic determinations for children. Aware of the limitations and potential of current guidelines, the proposed system aims to overcome the limitations of traditional guidelines by using artificial intelligence, which enhances adaptability to specific cases. Furthermore, it is envisaged that judges will always have the possibility to modify the decision proposed by the algorithm in order to adapt it to the specific case. The judge's discretion is not eliminated, similar to, for instance, with the aforementioned '*Tabelle di Milano*'. The proposed framework is only applicable if the judge shares the application's assumptions, thus maintaining a human-centered approach to justice. It is important to stress that the use of A.I.A.Co. is not intended to be mandatory for judges (also because the current state of the art does not allow it). Certainly, the Project can and should be improved and implemented to enhance the fairness and effectiveness of the legal system, but it seems to be a good starting point for those interested in approaching the issue of predictive justice.

[28] See Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139, and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act).

Annex

The interested reader can find in this section a sample of the defense deeds that have been sent to the Judges for labeling, see Section 'The A.I.A.Co. project', several figures produced with SHAP, and 3 examples of counterfactuals, see Q4 in Section 'Validation survey'. These figures can be used to understand from various dimensions why the A.I.A.Co. framework outputs a certain decision.

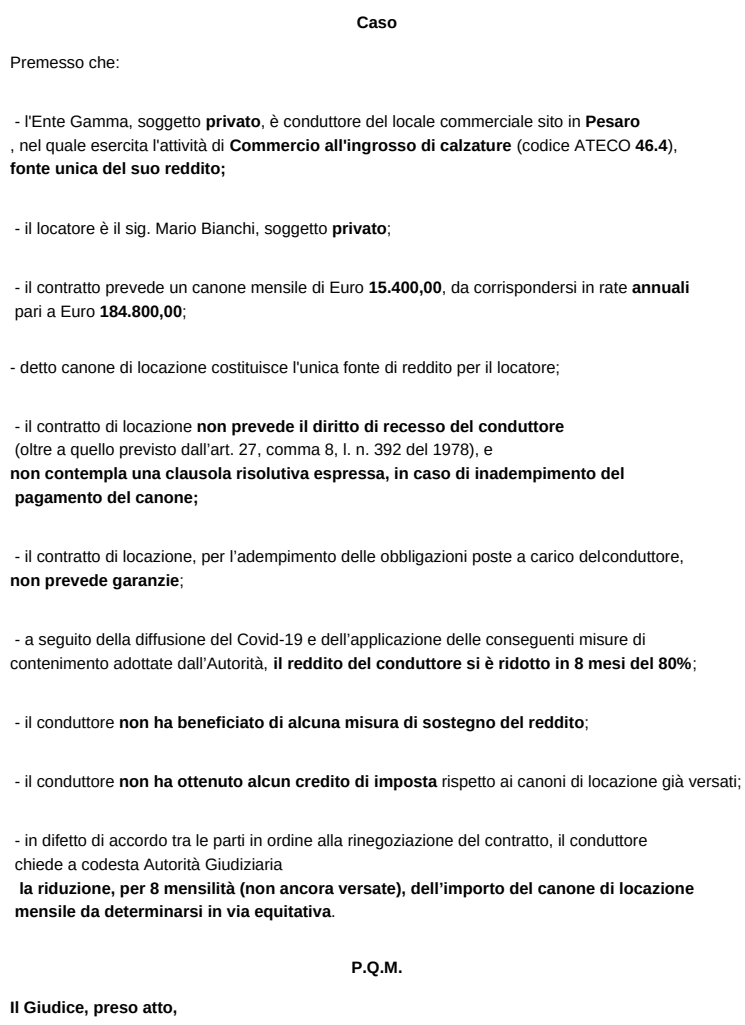

Caso

Premesso che:

- l'Ente Gamma, soggetto **privato**, è conduttore del locale commerciale sito in **Pesaro** , nel quale esercita l'attività di **Commercio all'ingrosso di calzature** (codice ATECO **46.4**), **fonte unica del suo reddito;**

- il locatore è il sig. Mario Bianchi, soggetto **privato**;

- il contratto prevede un canone mensile di Euro **15.400,00**, da corrispondersi in rate **annuali** pari a Euro **184.800,00**;

- detto canone di locazione costituisce l'unica fonte di reddito per il locatore;

- il contratto di locazione **non prevede il diritto di recesso del conduttore** (oltre a quello previsto dall'art. 27, comma 8, l. n. 392 del 1978), e **non contempla una clausola risolutiva espressa, in caso di inadempimento del pagamento del canone;**

- il contratto di locazione, per l'adempimento delle obbligazioni poste a carico delconduttore, **non prevede garanzie**;

- a seguito della diffusione del Covid-19 e dell'applicazione delle conseguenti misure di contenimento adottate dall'Autorità, **il reddito del conduttore si è ridotto in 8 mesi del 80%**;

- il conduttore **non ha beneficiato di alcuna misura di sostegno del reddito**;

- il conduttore **non ha ottenuto alcun credito di imposta** rispetto ai canoni di locazione già versati;

- in difetto di accordo tra le parti in ordine alla rinegoziazione del contratto, il conduttore chiede a codesta Autorità Giudiziaria **la riduzione, per 8 mensilità (non ancora versate), dell'importo del canone di locazione mensile da determinarsi in via equitativa**.

P.Q.M.

Il Giudice, preso atto,

Figure 5: *One of the artificial defense deeds sent to the Judges. The Law Unit selected a set of realistic features and likely constraints, together with weights, to attenuate the distribution shift from the training data to actual cases. Then, 600 data points have been sampled using the chosen feature distributions. Finally, each of these data points was used to produce an artificially generated defense deed in a human-readable PDF document in the Italian language. This is the first of these 600 defense deeds, as it was transmitted by email to a Judge.*

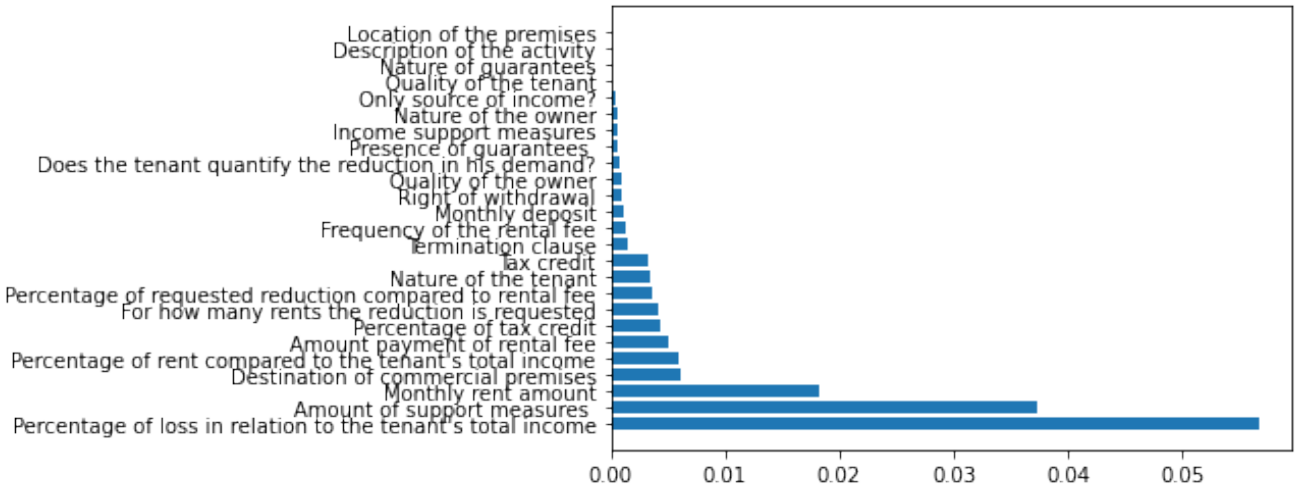

Figure 6: *The original 25 expert-selected features, and their corresponding SHAP values. We removed features whose SHAP value was less than* 10^{-5}*, and this leaves 21 features.*

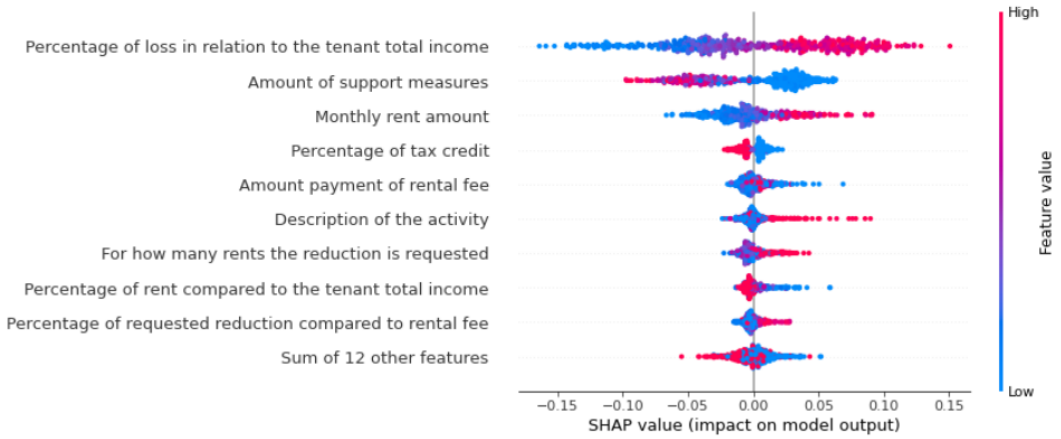

Figure 7: *Scatter plot with SHAP values of all instances for the most influential features. The values are grouped by the features on the y-axis. For each group, the color of the points is determined by the value of the same feature. The features are ordered by the mean SHAP values.*

Figure 8: *Dependence plot for the SHAP values of the feature 'Percentage of loss in relation to the tenant income', the most important feature.*

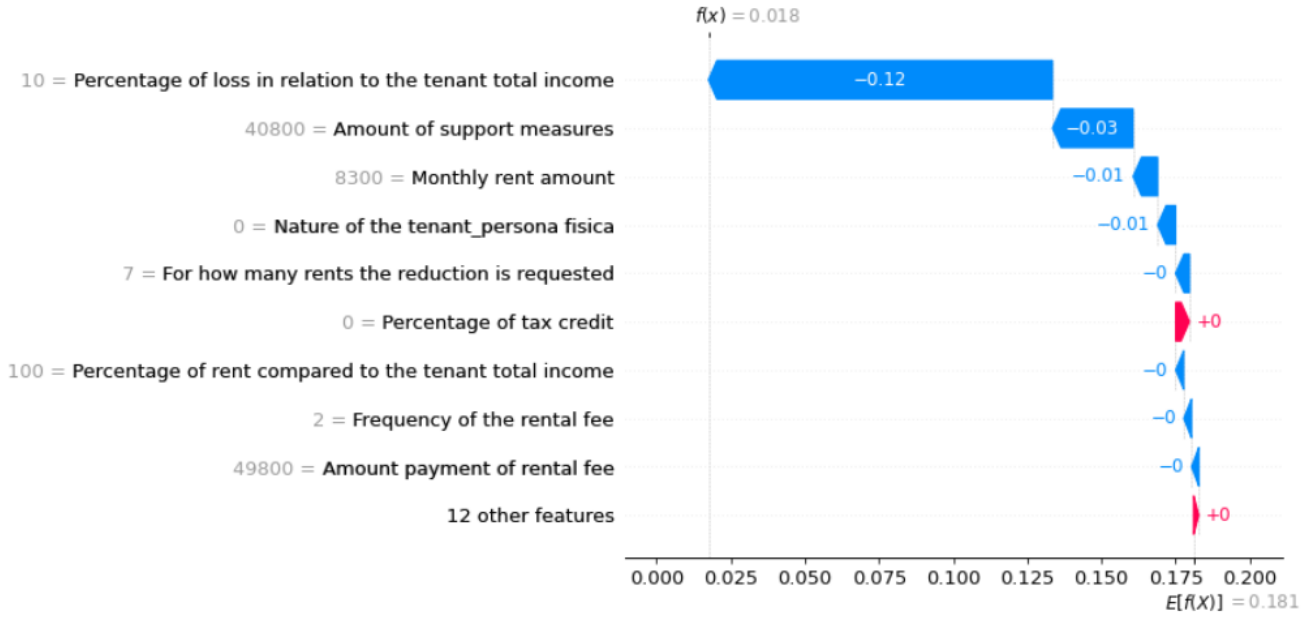

Figure 9: *Waterfall plot. SHAP values always sum up to the difference between the model outcome when all features are present and the model outcome with no features at all. Thus, SHAP values of all the input features will always sum up to the difference between the baseline (expected) model output and the current model output for the prediction being explained. The easiest way to see this is through a waterfall plot that starts at our background prior expectation* $\mathbb{E}[f(X)]$*, and then adds features until we reach the current model output* $[f(X)]$*.*

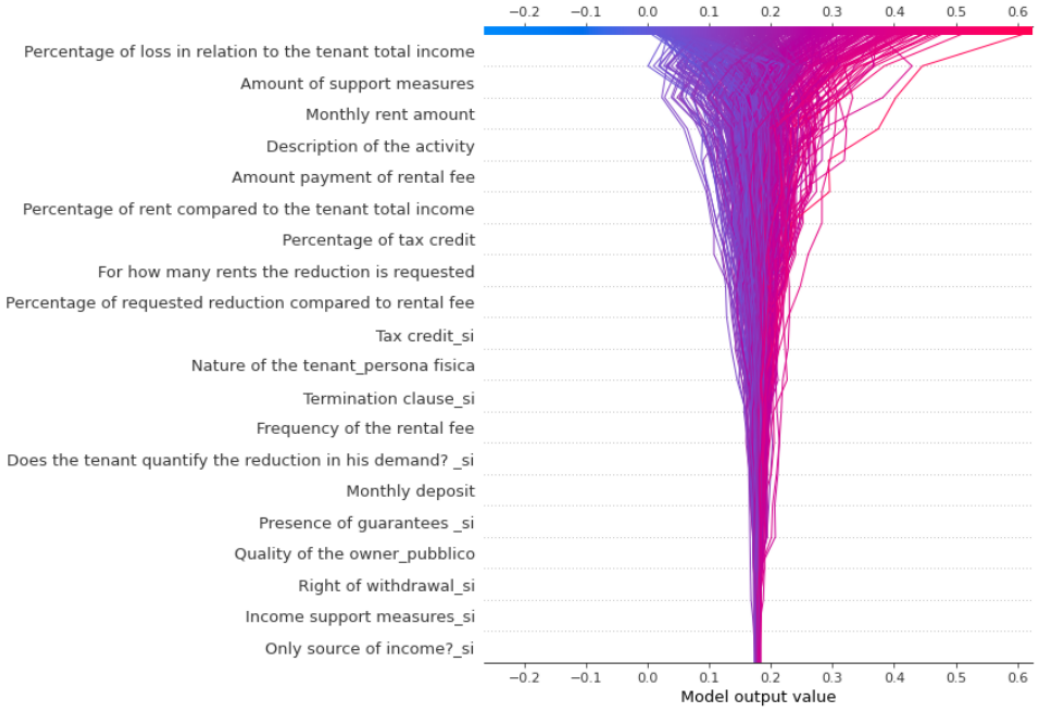

Figure 10: *To understand how our model makes predictions we need to aggregate the SHAP values. One way to do this is by using a decision plot. In this figure, each line corresponds to one of the model's decisions. It starts at the same base value and ends at its final predicted number. As you move up from each feature on the y-axis, the movement on the x-axis is given by the SHAP value for that feature. This gives you similar information to a waterfall plot.*

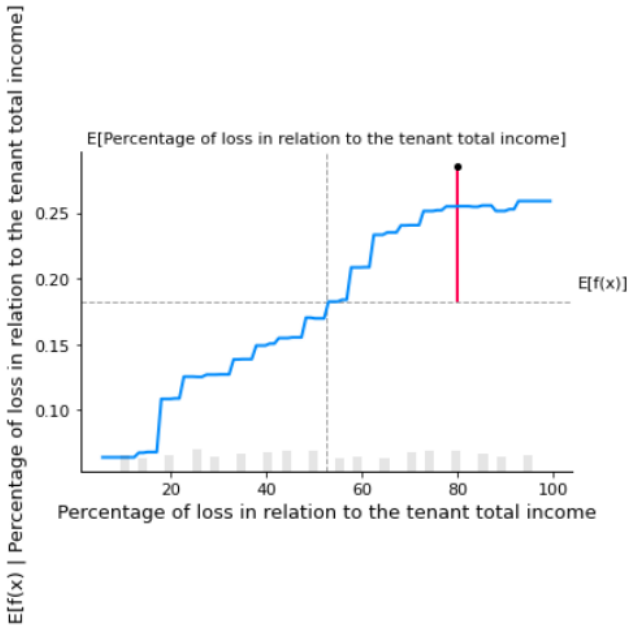

Figure 11: *Scatter additive plot for the most important feature.*

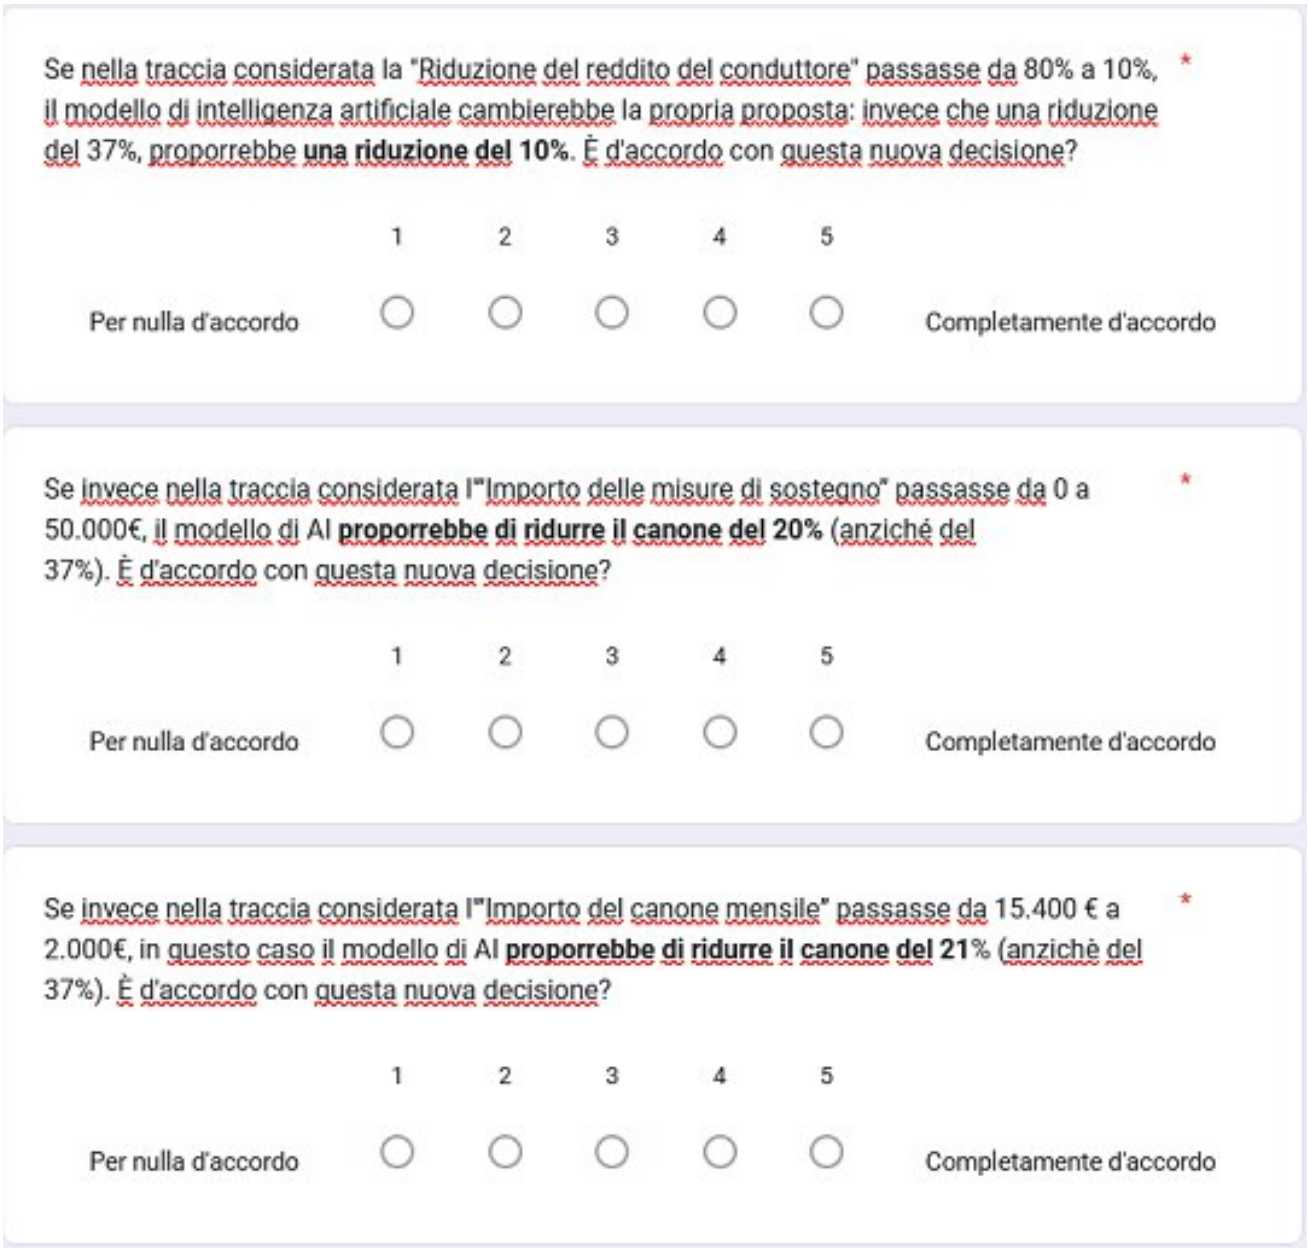

Figure 12: *Example of 3 counterfactuals proposed in the survey. The top one says: 'In this defence deed, if the feature [...] was 10% instead of 80%, the AI system would propose a reduction of 10% instead of 37%'. Participants where asked to rate those counterfactuals with a score from 1 to 5.*