We would like to thank the editor and the reviewers for their constructive comments. See below our point-by-point response.

**Reviewer 1**

**1. At lines 88-94, the authors suggest the model can 'be used to rapidly identify cases', perhaps in 'submitted cases'. It remains fundamentally unclear to me what text the authors imagine the model will be operating on to make ex ante predictions, given that it works solely on text in completed judgments. Is it not at least worth *identifying* what the text is that the model would operate on in order to make its 'rapid' ex ante predictions? Otherwise, the risk is that the argument will appear to be limited to this: 'Our model allows us to analyse how the facts are summarized in published cases and predict how, later in that very same published case, the case will be resolved'. This may be interesting in itself, but I'm not sure it's what the authors want the model to do.**

Our main argument in favour of ex ante predictions of outcomes rests on the premise that there is enough similarity between (at least certain) chunks of text of completed judgments by the Court and other kinds of textual materials, to wit: (a) applications lodged with the Court, (b) briefs submitted by parties with respect to pending cases making ECHR-related legal arguments and, possibly, (c) sections of domestic judgments that touch upon ECHR-related issues (whether an application to the Court has been made or not). Since our predictive NLP approach (and not the specific supervised model *per se*) seems to work reasonably well with text chucks from published judgments issued by the Court, it could also be further tested to assess whether it in fact can generalise to these other kinds of (in our view quite similar) texts. Unfortunately, this is something that we cannot test at the moment, because we do not have access to the data set that will allow us to do so (and we will not have access to such material in the foreseeable future, since the Court does not easily give access to lodged applications or briefs submitted by parties, which would be the most natural candidates). We thus used published judgments as proxies for the material to which we do not have access. At the very least, our work proves the following point (and is therefore at least a 'proof of concept' in this limited sense): *if* there is enough similarity between the chunks of text that we analysed and/or applications and briefs, as we believe there is, *then* NLP approaches can be fruitfully used to predict outcomes with a certain degree of reliability (at the very least, with a degree of reliability that appears to exceed by far the random 50% distribution). So, at the very least, our work provides some justification for further research, potentially on a wider scale and with the use of materials to which we do not at the present moment have access. We hasten to add that this kind of research programme requires different kinds of resources than the ones currently at our disposal, since it renders imperative a close cooperation with the Court itself, with all that such a cooperation entails. Moreover, we provided various reasons to lend support to our view according to which the sections of the cases we analysed can *indeed* be used as proxies for these other sorts of materials (i.e. applications and briefs), and we shall get back to these reasons in the remaining sections of our answer. Overall, we do not believe that a wholesale scepticism with regard to such uses of sections of cases is warranted. We have incorporated the above clarifications in the article.

**2. The new section at lines 140-166 seems problematic, unless I have misunderstood key aspects of it. First, the authors say that the ECtHR has very limited fact-finding powers (true). But the authors then move on, without citing any authority, to say that this means that 'in the vast majority of cases, [the ECtHR] will defer…to the judgments of domestic courts that have already heard and dismissed the applicants' complaint'. This is problematic in various ways (not least that it implies that the domestic courts hear and dismiss complaints on the same legal questions that the ECtHR does, which seems to suggest the ECtHR is an appellate court – more on this below). More problematically, the authors' logic is unclear: why wouldn't the ECtHR defer to the summary of the facts prepared by (e.g.) the government lawyers? Moreover, even if it were true that the Court defers in this manner 'in the vast majority of cases', surely it should not be difficult to find a range of law journal articles and analysis supporting this legal or procedural proposition? And finally, this logic would suggest that any predictive model should look principally or exclusively at the domestic courts' factual summaries, would it not?**

There are various things to note here. First, we explicitly say that deference to domestic judgments is to do *only with summaries of the factual background of the case* and not with anything else. The exact phrase we had used is this "in the vast majority of cases, it [the ECtHR] will defer, *when summarizing the factual background of a case*, to the judgments of domestic courts" (emphasis added). As a result, we nowhere either said explicitly or implied that domestic courts hear and dismiss complaints on the same legal questions as the ECtHR does. In any event, though, we have taken note of the reviewer's comments and have revised the text accordingly, to remove any possibility of misunderstanding. Second, we are of the view that the main thing to underline is that the facts of the case, in the vast majority of judgments where the Court does not use its powers of investigation, are *not in dispute by the parties*. Accordingly, for the vast majority of cases, the content of those facts has been fixed by the rules of procedure (and evidence) used at the domestic level by national jurisdictions. Moreover, in the vast majority of cases where the facts are thus fixed, the main question that the Court has to answer is a legal (whether there has been a violation of some ECHR-protected right) and not a factual (what are the facts of this particular case) one. We thus believe that it can be reasonably held that, in view of the above, the question of the exact source used to summarize the facts is of secondary importance. We also wish to stress that 'reasonably' in the above sense does not mean 'bulletproof' for all intents and purposes. As a matter of empirical fact, we can never be sure about the exact ways in which the Court arrives at summaries of the facts of each case and this is a limitation to our analysis that we have revised our text to take account of. In addition, we have revised the text of our article to provide references to academic work on how registry lawyers prepare summaries of the facts for judges (an issue on which some -limited- socio-legal research has been conducted). Again, our argument here depends on certain (in our view) reasonable assumptions. We cannot prove these assumptions, since we do not have access to the materials themselves, so readers are invited to read our argument as a hypothetical: *if* these assumptions are met (which we think they do), *then* certain consequences follow. Third, In order to back our claim that the Court mostly defers to domestic courts when determining the

factual background of a case, we provided one reference to the leading research on the matter and one more reference to a textbook by a leading legal practitioner of the ECHR.

**3. Second, the authors say that 'the Court cannot openly acknowledge any kind of bias on its part' and therefore 'on their face, summaries of facts…have to be at least framed in as neutral…'. The significance of this point is unexplained. Are the authors arguing that the Court does, in reality, prepare neutral summaries? Or that it may well be biased but that it must hide that bias? How does this help the argument about making ex ante predictions of any sort? Moreover, the authors do not seem to appreciate here that outright bias is only one problem: the bigger problem for their argument is the possibility of perfectly rational differential emphasis by the judges/registry/etc of facts that they know will be significant *because they are also involved in reaching the legal conclusions*. Unless I have misunderstood the ECtHR's procedure, in which case I apologize.**

As we have already mentioned, we are not in a position to either corroborate or refute empirically the proposition to the effect that the fact summaries prepared by the Court are neutral or not. The point we tried to make is merely that, since the Court has to (at least) appear unbiased to both parties in a dispute, it has an interest in presenting the facts of the case in what appears to the parties as being an unbiased way. There is thus an incentive that the facts are characterised in a way that does not (e.g.) hide certain crucial events or misrepresent them. So this, again, is an argument in support of the hypothesis that the chunks of text found in judgments can indeed be reasonably taken to stand as crude proxies for other kinds of texts (lodged applications/briefs of parties), i.e. to the main candidates for textual analysis of ex ante predictions of outcomes to which we do not have access. Moreover, we fully appreciate the possibility of forms of non-outright bias (and we have revised our text accordingly), but, as already stated at various points, we cannot control for the (eventual) presence of that variable, because we cannot compare the various versions of fact summaries due to lack of data. We have revised our text accordingly to make this clearer.

**4. Third, the absence of disputes before the ECtHR about the facts does not mean that there cannot be different facts emphasized or prioritized by the judges/registry in light of the analysis that they most likely know will follow.**

We agree with the reviewer, and we have revised our text accordingly (see lines 160-162 of the revised text).

**5. Fourth, the authors say 'the "Circumstances" subsection is the closest (even if sometimes crude) proxy we have to a reliable textual representative of the factual background of a case'. Perhaps this is so, but one may wonder about what 'reliable' is worth here without a clearer sense of the model's utility. Surely the facts as summarized in government and/or applicant arguments might be worth a look if the goal is to assist the Court/lawyers in making ex ante predictions about how cases will be decided, even if they are not 'reliable' in the way a peer-reviewed paper might be**

**reliable?**

The reviewer is right, but, as a matter of fact, and as we have already said, we do not have access to these texts and this is the main reason we have used such crude proxies. We submit once again (see our response to the first point) that such access requires other kinds of resources than those we currently have at our disposal. We also want to stress that, in any event, analysis of such further texts, whose obtention requires other kinds of resources, would be a logical step to take once we have *some* evidence, as our article suggests we do, to the effect that an NLP approach can indeed, in principle, be used to predict outcomes with a certain degree of reliability. In any event, we struck out the word 'reliable', which we do not think adds something substantial to our argument, given the caveats already in place.

**6. At several points (line 325, line 340), the ECtHR seems to be referred to as an appellate court. It is not an appellate court. This means, (1) the range of orders and remedies available to the ECtHR are not those of an appellate court, with consequences for its analysis; (2) the ECtHR will frequently be applying different *legal* tests to those applied by the domestic appellate courts (eg, 'was there a violation of Art5 or Art6?' vs 'was the conviction unsafe?'), and (3) therefore different *facts* or different emphasis on those facts may be of interest to the ECtHR than those that were of interest to the domestic court.**

We agree with the reviewer on the 'appellate court' point and we have revised our text accordingly. On the question of facts, see our response to the first point. It is of course possible that different facts may be of interest to the ECtHR (this is always the case when, e.g. the Court uses its fact-finding powers to investigate whether a Convention right was violated at the domestic level, even if such use is rare) but this just restates the point that the relevant sections we used are only crude and imperfect proxies for the facts of the case.

**General comment:** We agree that there are various constraints that curtail the utility of our work. We have quite clearly stated that we do not have access to lodged applications and briefs submitted by parties, nor will we have in the foreseeable future. Under these conditions, we proceeded to a crude emulation of the process, by using instead chunks of text contained in the judgments themselves, that we believe can reasonably stand as crude proxies for the real thing (summaries of facts and legal arguments in applications and briefs) to provide an initial test for an NLP approach. It is precisely, we think, the *success* of this initial test that corroborates the idea that further testing is needed, with recourse to other kinds of data. Again, we stress that access to such data would open avenues for further research and analysis. However that requires the active collaboration with the Court, and therefore a level of trust that exceeds by far what we can work with at the present moment.

**Reviewer 2**

**7. The authors commented during the review process that there are some barriers for accessing the data, from the EctHR portal. Besides, the authors mentioned that accessing cases from domestic courts is not straightforward either. I would like to see a comment regarding data access issues, perhaps a sentence or two in the conclusion section. Are data access issues stopping or slowing down emerging research as the one presented in this paper? Should cases in the EctHR and domestic cases be easily and freely available? If that it the case, is the EctHR making progress in open their repositories for public good?**

Data access issues slow down emerging research indeed. To the best of our knowledge the HUDOC database has not been designed to support large scale access. It is not possible to download the entire dump as in other large repositories such as Wikipedia. Of course making the data available easily accessible will further enable other studies considering the vast amounts of text and the rich associated metadata contained in HUDOC. We have added a couple of sentences in the conclusions section discussing the issue.