



# PEER REVIEW FRAMEWORK FOR PREDICTIVE ANALYTICS IN HUMANITARIAN RESPONSE

MAY 2021

THE CENTRE FOR  
HUMANITARIAN DATA



OCHA

centre for humdata



## 1. INTRODUCTION

Humanitarian decision-makers have called for the increased use of predictive analytics to anticipate and better respond to humanitarian crises. However, translating the outputs of predictive models into timely and appropriate responses remains a challenge for several reasons:

- First, there is no common standard for **documenting predictive models and their intended use**.
- Second, there is no common standard or mechanism for **assessing the technical rigor and operational readiness** of predictive models.
- Third, the development of predictive models is often led by technical specialists who may not consider important **ethical concerns** related to the application of models in humanitarian contexts.

Since 2018, the United Nations Office for the Coordination of Humanitarian Affairs (OCHA) Centre for Humanitarian Data ('the Centre') has been working with our partners to understand the state of model development and use in humanitarian operations. We have noted a clear desire for quality assurance of models by partners, with the Centre identified as having a unique role to facilitate a peer review process.

The following *Peer Review Framework for Predictive Analytics in Humanitarian Response* ('the Framework') aims to create standards and processes for the use of models in our sector. It is based on research with experts and stakeholders across a range of organizations that design and use predictive models. The Framework also draws on best practices from academia and the private sector. The purpose of applying the Framework is to offer advisory support in the development and implementation of predictive models. The results of the review do not constitute an endorsement on behalf of OCHA or the United Nations.

In addition to the work on quality assurance through peer review, the Centre's predictive analytics workstream validates and develops models for use in humanitarian operations. Learn more about these efforts on the Centre's website: <https://centre.humdata.org/predictive-analytics/>.

## 2. ROLES IN THE PEER REVIEW PROCESS

There are three main roles in the peer review process: the Client, the Reviewers, and the Moderator. Additional stakeholders (individuals or organisations who could be affected by the use of the model or who may use model outputs to take decisions) may be consulted during the review process.

The Centre leads the process and works with experts to complete the peer review. The Client is asked to identify a single focal point for the process, although different colleagues may need to be involved for each step.

The roles in the process are:

- The Client: an organization submitting a model for peer review.
- The Reviewers:
  - Technical Reviewer: an individual with demonstrated expertise in data science or statistics and their application in the relevant context.
  - Ethical Reviewer: an individual with demonstrated expertise in practical and humanitarian ethics.
- The Moderator: a member of the Centre's predictive analytics team designated by the Team Lead on a case-by-case basis. In addition, the Centre's data responsibility team will support the process as needed.

## 2.1 Reviewer Selection and Reviewer Pool


The Centre invites experts to submit an application to become a Reviewer in the technical or ethical domain. [A Call for Reviewers](#) will be periodically promoted by the Centre in order to attract and maintain a group of experts to support the peer review process.

Once accepted, the Reviewer will become part of a Reviewer Pool managed by the Centre. The Moderator will select Reviewers based on availability and a match of skills for the model in reference. Reviewers will not be assigned to review models submitted by their own organization or in case of any other potential conflict of interest. The Reviewer role is unpaid.

## 3. STEPS IN THE PEER REVIEW PROCESS

Partners interested in submitting a model for peer review are asked to submit a request through [this form](#). The Centre will engage with the Client for an initial meeting and process walkthrough. A Moderator is assigned to the review and agrees with the Client on a timeline for the process. The Moderator provides the [Model Card template](#) to the Client for the formal submission of the model.

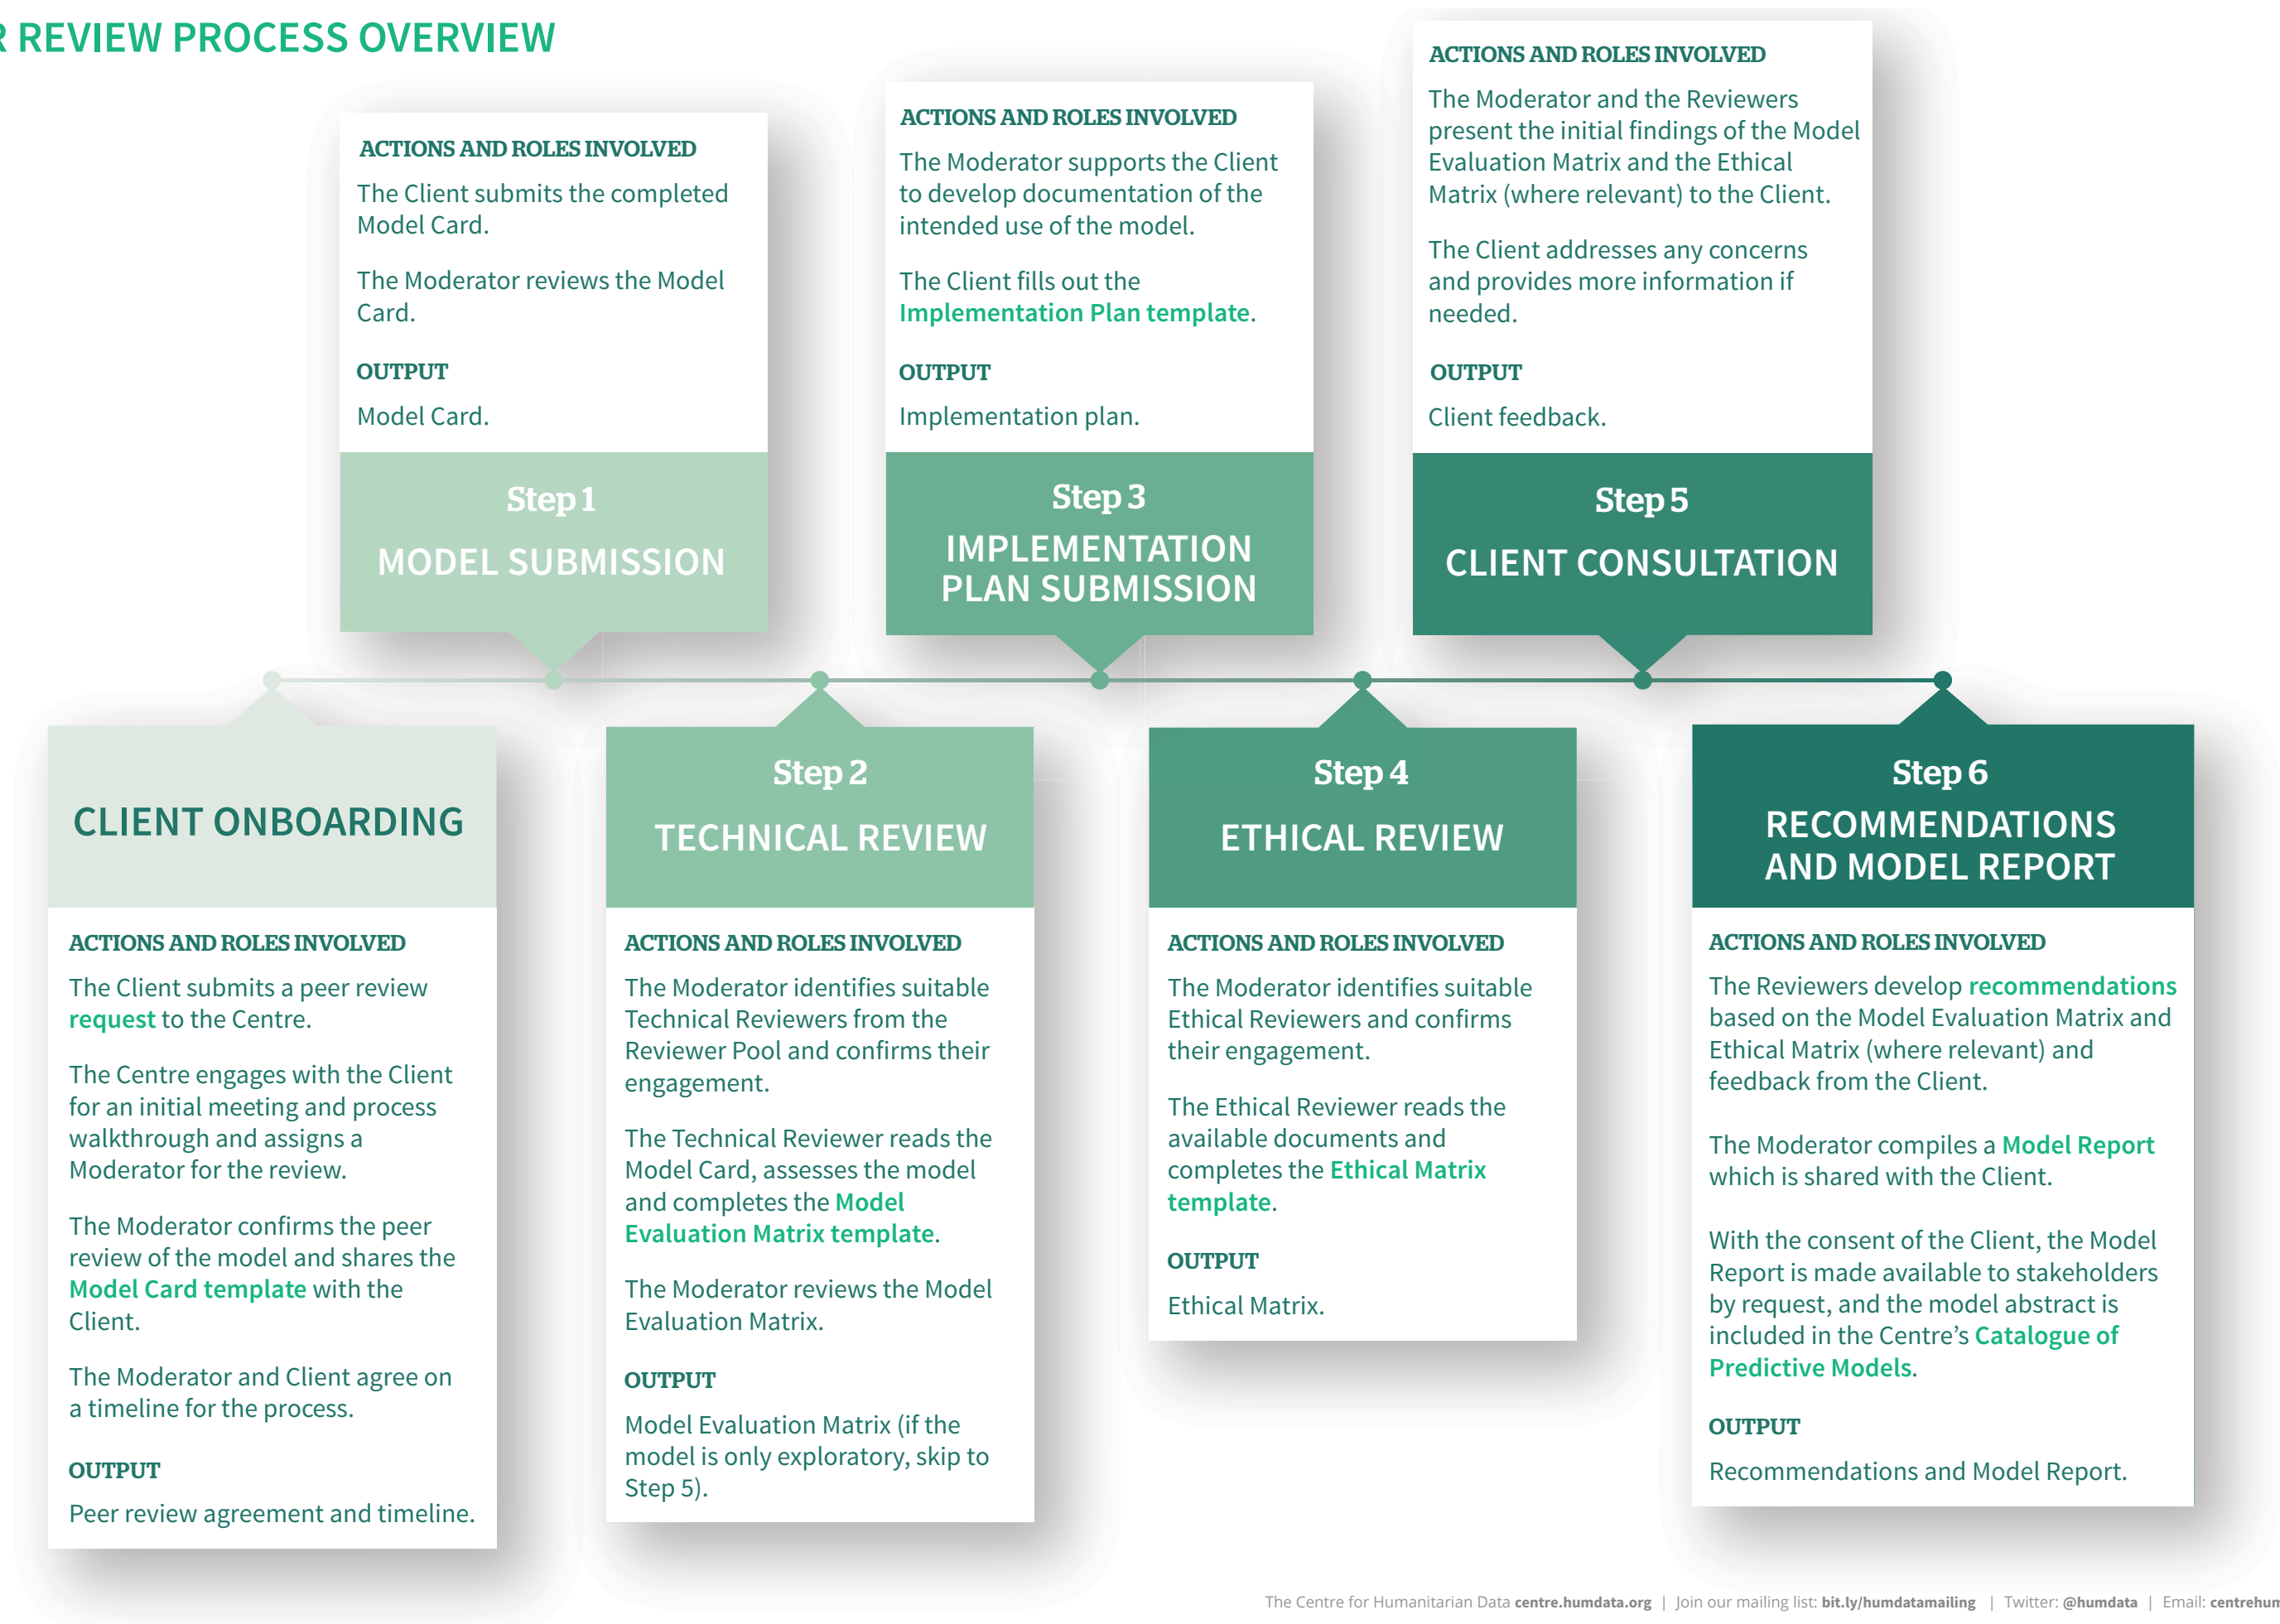
The process consists of six steps:



The Framework aims to support models that are either fully developed or exploratory. Models that are fully developed and have a defined or intended implementation plan will go through the complete peer review process. In the case of exploratory models, it is not possible to evaluate the ethical concerns as the actions that the model would generate are not yet defined. Review of these models will therefore skip steps 3 and 4.

The duration of the review process will depend on the steps that need to be completed and the quality of the submission by the Client. A complete review of a fully developed model should take around eight weeks.

# PEER REVIEW PROCESS OVERVIEW



## Step 01 MODEL SUBMISSION

In the first step, the Client submits the model for peer review by completing the [Model Card template](#), with support from the Moderator if needed. The Model Card includes the information required throughout the peer review process, including:

- The scope and the intended use of the model.
- The methodology of the model, including the reliability of input datasets, potential biases and the model assumptions.
- The evaluation process of the model and the reliability of the model output.
- The operational readiness of the model for humanitarian contexts.

Providing this information through the Model Card helps the Centre to standardize documentation, allowing stakeholders to compare candidate models.

## Step 02 TECHNICAL REVIEW

In the second step, the model undergoes a technical review based on the information in the Model Card. The output of this step is a completed [Model Evaluation Matrix](#) which highlights the technical strengths and weaknesses for the model's application in humanitarian action.

The Technical Reviewer will assess strengths and weaknesses across the following sections of the Model Card:

- Intended Use
  - Clarity in the definition of use-cases.
  - Evaluate the understanding of the model's boundaries, limitations and constraints.
- Model Development
  - Quality and representativeness of model's input datasets.
  - Robustness of model methodology.
- Model Evaluation
  - Reliability of projections with respect to existing benchmarks.
  - Understanding of model failures.
- Operational Readiness
  - Usability and timeliness of model.
  - Clarity in the definition of roles and responsibilities.

If the model being reviewed is only exploratory and has no intended use defined yet, the Moderator will skip steps 3 and 4 and move to set up the consultation with the Client. For models that are fully developed and with an intended use already defined, the Moderator will continue with step 3 in the process.

## Step 03 IMPLEMENTATION PLAN SUBMISSION

In the third step, the Moderator supports the Client to develop documentation of the intended use of the model. The output of this step is a completed [Implementation Plan](#) that summarizes the concrete actions that the model will trigger or inform and clarifies the roles and responsibilities for implementation of the model in humanitarian action.

## Step 04 ETHICAL REVIEW

In the fourth step, the Ethical Reviewer identifies key issues associated with the model and its intended use, assesses the related ethical concerns of different stakeholders, and offers recommendations for how to address these concerns. The output of this step is a completed **Ethical Matrix**.<sup>1</sup> During this step, the Ethical Reviewer will complete a draft matrix, request further information and feedback from the Client, and consult with the Moderator and colleagues from the Centre's data responsibility team as needed. A detailed description of the process for completing the Ethical Matrix is provided in Annex A.

## Step 05 CLIENT CONSULTATION

The Moderator will convene a call with the Client and the Reviewers to present and discuss the initial findings of the review. The Client will have the opportunity to provide feedback and to complement the information provided in the Model Card and the Implementation Plan. The Client may also explain if and how any weaknesses in the model have been addressed. If the Client reports that they have taken actions to improve the model, those measures will be reflected in the final versions of the Model Evaluation Matrix and Ethical Matrix.

## Step 06 RECOMMENDATIONS AND MODEL REPORT

In the final step, the Reviewers develop **recommendations** based on their review of the model and the feedback from the Client. When developing recommendations, the Technical and Ethical Reviewers assess additional areas which need further attention to mitigate potential concerns in the implementation of the model.

The Moderator produces a **Model Report** based on the information in the Model Card, Model Evaluation Matrix, Implementation Plan, the Ethical Matrix, the Client consultation and the Reviewers' recommendations. The Moderator will share the Model Report with the Client. The Client may decide whether the Model Report should be made publicly available to stakeholders, on request or be kept confidential. An abstract of the model will be included in the Centre's **Catalogue of Predictive Models**.

---

## FEEDBACK

The Centre invites individuals and organizations working in the humanitarian, academic, research and private sectors to engage with us on the peer review process. Please send feedback on the Framework to [centrehumdata@un.org](mailto:centrehumdata@un.org).

## ACKNOWLEDGEMENTS

The initial draft Peer Review Framework was developed by Dani Poole who worked as a Data Fellow with the Centre during the 2019 Data Fellows Programme in June and July The Hague. As part of her research, Dani conducted interviews with over 20 experts including data scientists, researchers, ethicists, and decision makers spanning the humanitarian, academic, and private sectors. The draft Framework was further detailed and refined with input from Leonardo Milano, Evan Tachovsky, Madeline Lisaius, Stuart Campo, Jos Berens, Sarah Telford, Eleonore Fournier-Tombs, Manu Singh, Kirsten Gelsdorf, Nicholas Bodanac, and Fanny Weicherding, among others. We appreciate the time and consideration of the many people who have contributed to this process and will continue to update the Framework on a regular basis.

<sup>1</sup> The Ethical Matrix has been adapted from the approach developed by Cathy O'Neil. For more information, see O'Neil, C. Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown Publishing Group New York, NY, USA, 2016.

## ANNEX A

### ETHICAL REVIEW GUIDE

This Ethical Review Guide describes the process for completing the Ethical Matrix. Use the [Ethical Matrix template](#) to record your observations as you move through the process.

#### A. List parameters for the development and implementation of the model.

*What are the parameters limiting the model?*

---

List all documents containing parameters for the development and implementation of the model. This may include applicable legal frameworks, organizational procedures, sectoral or organizational codes of ethics and other guidance or governance documents. The Implementation Plan should include a list of key documents for consideration as part of your review.

Consider the following types of parameters:

- Legal, e.g. relevant clauses in applicable data protection legislation. Who would be at risk if the algorithm fails?
- Contractual, e.g. relevant clauses in contracts applicable to the different stages of the model development and use.
- Procedural, e.g. relevant sections in applicable data protection and privacy policies.
- Ethical, e.g. relevant principles and related obligations in (sector-specific) ethical guidance, codes of conduct and similar documents.

#### B. Identify the stakeholders in the context of the model and its intended use.

*Who is affected by the intended use of the model?*

---

The Client will have provided their understanding of relevant stakeholders in the Implementation Plan, but the Ethical Reviewer should consider whether any additional stakeholders should be included in the Ethical Matrix. Stakeholders are actors who could be affected by the use of the model. To identify who is a stakeholder, consider:

- Who is responsible for implementing the actions based on model output?
- Who benefits from this model when it works as intended?
- Who would be at risk if the model fails?
- Who is directly involved in the development and implementation of the model?

Stakeholders may include individuals, groups of individuals (e.g. specific demographics or populations), organizations or specific teams within them (e.g. data team, program delivery team), or government agencies. Insert the identified stakeholders in the 'Stakeholders' tab in the Ethical Matrix template, listing them in order of priority. Determine the order of priority based on the centrality of protection, the humanitarian principles and the context of model implementation.

The below list contains examples of stakeholders that should typically be included. These stakeholders could also be split into subgroups if needed.

- **Affected population**
- **Modeller**
- **Senior decision-maker**
- **Host government**
- **Donor**
- **Local NGOs / response organizations**
- **International NGOs / response organizations**
- **Model reviewers**

### C. Identify potential issues.

*What issues may arise in the implementation of the model?*

---

Based on the content in the Model Card, the Model Evaluation Matrix and the Implementation Plan, generate a list of issues that may arise in the implementation of the model. Insert the identified issues in the top row of the Ethical Matrix. Issues should be listed in order of likelihood, placing the issue most likely to occur at the left of the list.

A non-comprehensive list of potential issues includes:

- **Inaccuracy.** The model's output is inaccurate. *A cyclone impact model estimates that 100,000 people are severely affected by the event but the true number is 180,000 people.*
- **Inaction.** No action is taken based on the model output. *A model predicts a migration of people and no preparations are made in the destination area.*
- **Overreaction.** A decision informed by the model output leads to action that is too impactful compared to the need that should be addressed. *A model predicts a rise in market prices and aid agencies respond by oversupplying the area with foodstuffs, disrupting local markets further.*
- **False positive.** The model produces a false alarm. *A food insecurity model wrongly estimates that a district will face a food security crisis in the next 3 months.*
- **False negative.** The model misses an event. *A food insecurity model does not catch an upcoming food security crisis.*
- **Lack of transparency.** The algorithms are not made available or are not transparent. *The model has been developed by a private company which uses a proprietary software that is not made available.*
- **Gaming.** Stakeholders can act according to the algorithm's rules to achieve an intended outcome. *An organization systematically inflates the severity of needs of the beneficiaries they are assisting to have greater humanitarian financing.*



- **Statistical bias.** The data used to train the algorithm does not reflect the reality on the ground. *The displacement data available to train a displacement model only covers 7 out of 10 districts because of limited humanitarian access.*
- **Systematic bias.** The datasets used to train the algorithm do not reflect the full complexity of the reality. *A drought impact model doesn't take into account factors related to ongoing conflict and insecurity but only environmental factors in a country with ongoing conflict.*
- **Corruption.** Offering goods/services to make the algorithm present a more favorable outcome. *A private company developing an impact model for heat waves is offered compensation from a service provider to tune the parameters of the model to increase the probability of triggering financing.*
- **Politics.** The outcomes of the model are not acted upon due to competing political interests. *A model is developed without an agreed response plan and local actors are not involved in the implementation.*
- **Missing attributes.** The data used to train and validate the model has missing attributes. *The input data for population is not disaggregated by age and the number of children in the affected areas is obtained by extrapolating from the national average.*
- **Insufficient data.** Historical data on previous impacts is not available for proper testing of the model. *An epidemic model is only verified on one historical outbreak.*
- **Lack of consent.** The model doesn't align with the consent provided at data collection. *The model uses population density data from a private data provider released under a non-disclosure agreement.*
- **Privacy infringement.** The model allows individuals or vulnerable groups to be re-identified. *A flood impact model uses microdata on displacement population allowing identification of vulnerable groups.*
- **Lack of data security.** The model's inputs and outputs are accessed by a non-authorized entity. *The model is using sensitive population data that is stored in a non-secured data warehouse.*
- **Ossification.** The algorithm learns the current state of the world and predicts according to it. *Conflict and insecurity escalated in a country where a drought food security model is implemented but the model was tuned on historical data and only considers environmental factors.*

#### D. Identify specific ethical concerns of different stakeholders.

*What ethical concerns do the issues raise for different stakeholders?*

For each potential issue (columns in the Ethical Matrix), go row-by-row and consider the implications for each stakeholder group. Ask: If this issue occurred (e.g., if the model generated false negatives), would this create an ethical concern for this particular stakeholder group? If yes, how? Record these concerns in the relevant cell (stakeholder/issue) in the Ethical Matrix template provided.

Some stakeholders might need to be split into subgroups if those different subgroups have different concerns. You can add rows to the Ethical Matrix accordingly. Note that not all issues would affect all stakeholders, which should lead to 'N/A' in the Matrix.

This step is complete when the Ethical Reviewer believes that all relevant stakeholders and concerns are represented in the Matrix.

## E. Color-code concerns according to priority level.

*What's the priority of a given concern for a specific stakeholder?*

Review the concerns in the Ethical Matrix one-by-one and assign a priority level.

<b>N/A</b> the given issue raises no ethical concerns for this stakeholder.
<b>Medium Priority</b> the given issue raises ethical concerns that warrant some attention for this stakeholder.
<b>High Priority</b> the given issue raises ethical concerns for this stakeholder that warrant immediate attention and could compromise the use of the model.

### Priority Levels

Judging the level of priority can require a variety of sub-steps, including but not limited to:

- Considering independently the level of attention the various concerns should receive and the likelihood that the related issue will materialize for the relevant stakeholder.
- Seeking out complementary information to assess risk, such as probability estimates of model errors (e.g. false negative rate). These estimates will be available to the Ethical Reviewer in the Model Card.
- Assessing how to rank and weigh the intended use of the algorithm.
- Using data analysis or other research to validate the priority that the concerns should receive.
- Deciding how the different issues and related, potentially competing concerns of the different stakeholders will be balanced.
- Recognizing hard legal and procedural constraints as noted in the Implementation Plan.
- Considering the potential consequences of using the model versus not using it, i.e. of continuing the process that the algorithm would replace, including the possibility that there is no such process in place.

## F. Share draft Ethical Matrix with Moderator and participate in client consultation.

The Ethical Reviewer should share the draft Ethical Matrix with the Moderator when ready. The Moderator will convene a call with the Client and the Reviewers to present and discuss the initial findings of the review. During this call, the Ethical Reviewer will briefly explain their key observations as reflected in the draft Ethical Matrix. The Client will have the opportunity to provide feedback and to complement the information provided in the Model Card and the Implementation Plan. The Client may also explain if and how any weaknesses in the model have been addressed. If the Client reports that they have taken actions to improve the model, those measures will be reflected in the final versions of the Model Evaluation Matrix and Ethical Matrix.

## G. Finalize the Ethical Matrix and develop recommendations to address high-priority concerns.

The ultimate goal of the Ethical Review is to have visibility of all concerns and to address all high priority concerns in the model implementation. This means that none of the issues can impact a stakeholder in such a way that it would impair or cripple the use of the model, or that there is at least an ongoing monitoring system to alert relevant stakeholders of any concerns.

If the occurrence of an issue cannot be prevented, recommendations should focus on distributing the impact of the issue equitably so that it only impacts lower priority stakeholders.

In this final step, the Ethical Reviewer develops recommendations for each high-priority concern. This typically means developing a recommendation for any concerns in the top left hand 5 x 5 quadrant of the matrix. Recommendations should be included in the final report for the client. Recommendations could include:

- Updating the model itself (e.g. redefining the objective function, using different training and/or validation data).
- (Re)defining how the output of the model will be used to inform decisions or resource allocation.
- Other changes to the programme or service of which the model is part.
- Experiments, simulations or other research to gauge the impact of design changes on concerns.
- Developing automated monitoring systems and tools to continuously track the model performance in relation to key concerns.
- Specifying purposes or settings for which the model can and cannot be used.
- Explaining the model to users, affected people or other stakeholders.
- Engaging with stakeholder groups so they can help decide how their concerns will be addressed.

The list above makes clear that not all High Priority cells will be ‘resolved’ quickly or finally. For some concerns, the best we can do is create robust monitors and stay vigilant. In any case, new concerns will evolve as community values and constraints change over time. For these reasons, the outcomes of the Ethical Review should serve as a basis for regular revisions of the Ethical Matrix by the client, as the use of the model evolves.