Crowd Ethics

1. Idea

The idea is to use crowd intelligence to build the dataset of human ethics behaviors.

Considering we want to build a collection of human situations ethically accepted, we can start by asking an expert in ethics or moral philosophy to gather together recommendations and concepts of right and wrong conduct. This approach would be a good start to build an initial set of ethical situations; however, at a certain point we'll end up notice the limitations of this approach.

The fact is everyday life shows a wide variety of situations not covered by science where an ethical response is needed. In these cases, we rely mostly on our intuition and lifetime experience to decide between right and wrong. Each person holds a collection of lifetime experiences, some persons having a larger collection, others a smaller collection. The notion of right and wrong may have different nuances from person to person based on their accumulated life experience and the concepts built throughout their lifetime. We generally accept in our societies that having a richer life experience may enable someone to judge or to analyze a wider variety of lifetime situations from the perspective of right of wrong. However, this is not true in all cases, as it also requires the person not being biased in any way towards wrong doing, even if having a rich lifetime experience.

In conclusion, we can say that a single person or a limited group of persons cannot effectively build a comprehensive collection of ethically accepted situations, either due to the rich extend of lifetime situations or the potential bias regarding right and wrong concepts.

Therefore, we propose a solution involving crowd intelligence. For this purpose, the system will be exposed to a large number of persons, thus accessing the collective "wisdom" over a wide range of lifetime situations.

2. Building Dataset

Machines are generally able to learn from examples, provided in the form of datasets.

There are examples of machine learning technologies, like Google DeepMind, where starting from simple observations of games played by amateurs, the system has improved itself (through reinforcement learning) to reach a level so high that it was unbeatable for the Go World Champion.

Regarding machine ethics, we want to teach machines to distinguish between right and wrong, depending on situation encountered. The machine should be exposed to both right and wrong concepts, so to allow identifying the correct ethical behavior when confronted with a real-life situation. If a system doesn't know the notion of "wrong", we end up considering everything to be right (or ethical).

Once both concepts are efficiently trained, we need to make sure machines are biased towards the right ethical behavior when acting in a specific real-life situation. As, even if machines learn to differentiate right from wrong, this knowledge doesn't necessarily dictate they will use the right ethical behavior over time. There are situations where machines can "go bad" if there is a bias for other factors (like a short-term distraction). Thus, the learning process is essential, but behavior enforcement is also critical.

To build the dataset of ethical behaviors, we want to use human intelligence and intuition to classify real-life situations from ethical point of view.

We need a rich collection of real-life situations to a build a comprehensive dataset. Thinking on how to capture these real-life moments, the most accessible way nowadays is to use pictures. There are many websites featuring free stock photos, where amateur photographers capture lifetime moments and post them online to use for free. Being a professional photographer is not a mandatory feature, as we simply need lifetime snapshots, not directed nor prepared in any way.

When checking for free stock photos, we need to filter pictures for terms like "everyday life" or "people". This is our target of interest. One example of such a website is

https://www.shutterstock.com/search/everyday+life

Once we have gathered a large collection of pictures, we need to expose each picture to a large audience (or crowd) for classification. People will be requested to describe each picture in one sentence, identifying the main situation presented in the picture. Additionally, people will be requested to mark if situation presented in photo is ethical or not. There are cases where pictures don't show an ethical dilemma (i.e. where we need to decide between right and wrong). Example: a picture showing empty shopping trolleys. These examples can be ignored/skipped as they don't have a real value in the learning process.

To summarize, we propose the following process for building the dataset:

1. Identify websites/sources of free stock photos having points of interest like "everyday life" or "people" (other terms possible).

2. Build a large collection of images showing lifetime moments.

3. Use a crowd/mass collaboration platform to classify images from ethical point of view.

4. Collect all ethical behaviors into a machine learning dataset.

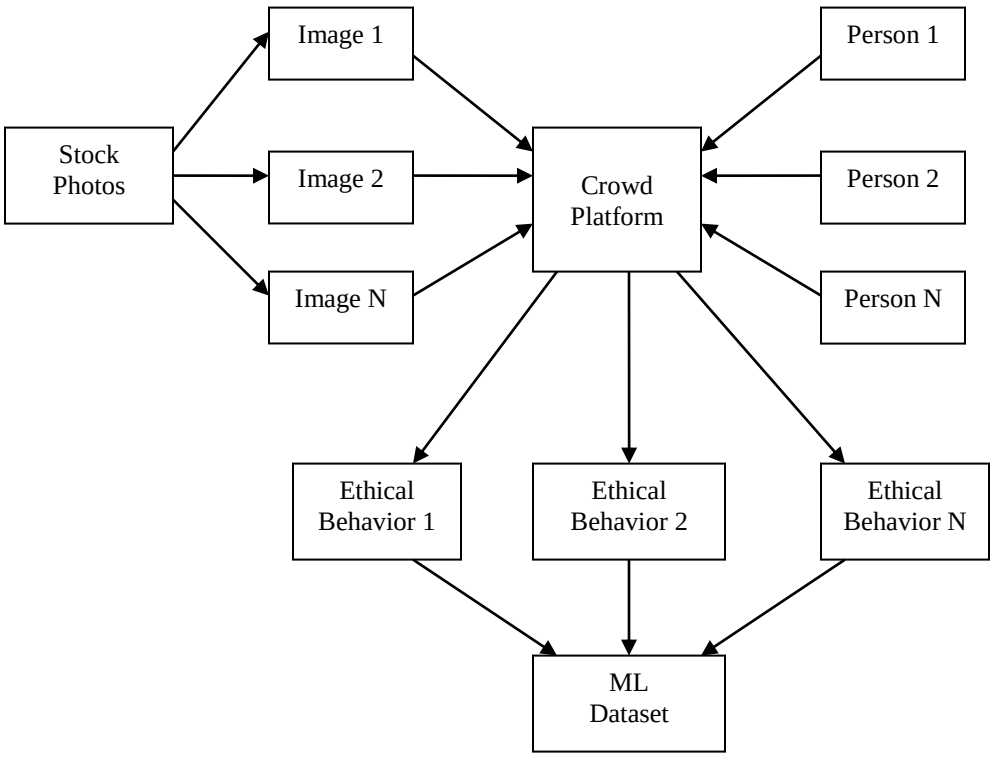

Figure. Building Dataset

3. Examples

Image 1

Description: Man washing dirty hands
Ethical: Yes

Image 2

Description: People minding distance on station platform
Ethical: Yes

Image 3

Description: Child doodles bathroom walls
Ethical: No

Image 4

Description: Child brushes teeth before going to bed
Ethical: Yes

Image 5

Description: Women washing vegetables before cooking
Ethical: Yes

Image 6

Description: Helping disabled woman get dressed
Ethical: Yes

Image 7

Description: Man steals wallet from woman's purse
Ethical: No

Image 8

Description: Throwing garbage on road out of car's window
Ethical: No

Image 9

Description: Pedestrian crossing street on red light
Ethical: No

Image 10

Description: Man punishing a dog with a bat
Ethical: No

4. Mass Collaboration Platform

The mass collaboration platform should provide the system functionality for organizing photos and allowing crowd access for classification.

Functionality-wise, the platform should be split in two areas:

1. Administrator Area

This section allows various functions such as:

- management of pictures (e.g. upload, updates)
- management of users (in case authentication on platform needed)
- generation of dataset for machine learning
- statistics, reports

2. Crowd Area

This section displays pictures for classification. Access to this area can be authenticated or not (depending on the purpose and scope of the project).

To make things more exciting, on every platform access there will be a new set of pictures randomly selected for classification.

Pictures selected can be either unclassified or already classified (so to have multiple opinions for each picture). Pictures having multiple classifications will require a general crowd agreement (from ethical point of view), so to be included in the learning dataset. It is actually preferable to have a crowd agreement on an ethical situation, so to avoid the bias of having a single person classification.

The final dataset will contain all ethical behaviors having a general crowd agreement.

The platform can consider strategies to motivate people advance the classification process. Such strategies can include gamification (like leaderboard, points, badges) or incentives (like prizes).