

How to Integrate Diverse Sources of Evidence

Learnings from the Lens Reputation System

Pedro Alcocer

Staff Data Scientist, Aave Companies / Lens Protocol

How do I integrate diverse sources of reputational evidence?



Represent your beliefs
about reputation signals as
beta distributions and
use Bayesian updating to
combine them.



What not to do:

$$\begin{array}{c} W1 \\ W2 \\ W3 \end{array} \times \begin{array}{ccc} S1 & S2 & S3 \end{array} = \text{Reputation}$$



What not to do:

$$W1 \times S1 + W2 \times S2 + W3 \times S3 = \text{Reputation}$$



Information about the
variance of your signals
is lost.



Variance is how the **quality** of a signal is represented.




You're throwing away
information about quality.



Weighting without variance



Weighting with variance



Weighting with variance




Signal 1



Signal 2



The **beta distribution** is bound between 0 and 1 and takes two parameters: **alpha** and **beta**.



Holds rAave POAP (high quality)

Beta(999, 1) \approx 1 out of every 1000 is a sybil

Has ENS address (medium quality)

Beta(99, 1) \approx 1 out of every 100 is a sybil

Has Twitter account (low quality)

Beta(9, 1) \approx 1 out of every 10 is a sybil



Bayesian updating is a way
of combining beliefs.



It's **really easy** to update
the beta distribution.
Just add alphas and betas.



$$\begin{aligned} & \text{Beta}(999, 1) \\ & \quad + \\ & \text{Beta}(99, 1) \\ & \quad + \\ & \text{Beta}(9, 1) \\ & \quad = \\ & \text{Beta}(1107, 3) \end{aligned}$$



Beta(3, 1) x 50 signals

=

Beta(150, 50)



Who would you rather trust?

Reputation 1



Reputation 2



Represent your beliefs
about reputation signals as
beta distributions and
use Bayesian updating to
combine them.



Thanks!

Pedro Alcocer

Leaf apps: pealco.lens

Bird app: [@pealco](https://twitter.com/pealco)

TG: [pealco_xyz](https://t.me/pealco_xyz)

