

Filed Under Seal

EXHIBIT D

>
> [REDACTED] is our on-call for the next two hours (and happens to be sitting right next to me :)), so I've given her a heads-up in case anything is needed.

[REDACTED] (5/22/2023 04:11:39 PDT):
>thank you!

[REDACTED] (5/22/2023 04:11:59 PDT):
>wouldnt the image alone be delete-worthy though, even without the caption? because it shows materials related to SSI?

[REDACTED] (5/22/2023 04:13:35 PDT):
>The razor blade is shared in an ED context _because_ of the account name ("selfharm_diaries") - which then makes this 'When there is a means of self-injury, even if passively depicted' under our graphic admission definition.

[REDACTED] (5/22/2023 04:13:58 PDT):

shared: sticker.png

[REDACTED] (5/22/2023 04:17:09 PDT):
>* SSI context - not ED (thanks for catching that [REDACTED]) @silent

[REDACTED] (5/22/2023 04:19:34 PDT):
>I thanks for handling this [REDACTED] Fatima, [REDACTED] [REDACTED] - If you were to prioritize the legitimacy risks we face, where would you place old content (pre our classifiers) compared to other risks like search and hashtags? @ [REDACTED] - I'm interested in the ops angle on this as well.

[REDACTED] (5/22/2023 04:21:19 PDT):
>@silent Hey [REDACTED] PREsc on-call currently seems to be dealing with it from the Ops side (@ [REDACTED] to confirm)

[REDACTED] (5/22/2023 05:24:23 PDT):
>Tbh it's really hard to prioritise because they all contribute to comms risk. I also imagine they may be linked - ie some accounts like these are coming up higher in search because we don't have classifiers running on old content, so they're not contributing to account score?

[REDACTED] (5/22/2023 05:24:50 PDT):
>Reporters are most likely finding these accounts / content by simply spending a couple of mins searching

[REDACTED] (5/22/2023 05:25:17 PDT):
>And if classifiers were running on old content I imagine we'd find and remove lots of content/accounts, meaning less for journalists and people in general to find

[REDACTED] (5/22/2023 05:25:36 PDT):
>(With the huge caveat that I'm not an ops or product expert!)

[REDACTED] (5/22/2023 05:32:18 PDT):
>#silent
>This is prohibitively expensive to do for all old content. @ [REDACTED] - what is the feasibility of an inexpensive option to do a level 1 scan of all old content to funnel it into more precise and expensive classifiers? My take on this is that we already do this kind of funneling for processing content live and the volume of old content is likely too massive for even this to be an option; I defer to your opinion on this.
>#remindus 9am

[REDACTED] (5/22/2023 08:08:47 PDT):
>#silent Is it possible to get the search terms used? That might help reduce the volume of old content we may have to run the model on. And how many of the searches were on complex entities?

[REDACTED] (5/22/2023 08:11:21 PDT):
>We can ask the journalist but I assume they'll have just been searching for obvious terms like 'self harm'

[REDACTED] (5/22/2023 08:11:57 PDT):
>And what do you mean re: searches being on complex entities? Don't you just search generally and then switch between tabs to see accounts / content / hashtags?

[REDACTED] (5/22/2023 08:13:19 PDT):

>@ [REDACTED] - as a rough way to scope the problem, the empty SeRP list and trigger term list can give you the search terms to target.

>The searches are usually not on entities. They are generic, the product classifies results by accounts, reels, tags etc

[REDACTED] (5/22/2023 08:14:17 PDT):

>I see got it

[REDACTED] (5/22/2023 08:16:12 PDT):

>You will find links to the lists I mentioned in this deeplink:

<https://docs.google.com/document/d/1G4ufiBdTzd0SDhiuvZpJgMnRjTcU3MegMIWJBwwwT5k/edit#bookmark=id.wpyac1qxzs57>

[REDACTED] (5/22/2023 08:41:38 PDT):

> [REDACTED] has this story picked up any team beyond Mail?

[REDACTED] (5/22/2023 08:42:07 PDT):

>no because the Mail hasn't published yet

[REDACTED] (5/22/2023 09:33:09 PDT):

>@ [REDACTED] - we have to line up capacity to handle appeals when we launch PV demotions for SSI & ED in the EU next year. Please add this item to your tracker with an assigned POC.

[REDACTED] (5/22/2023 10:05:15 PDT):

>hey @ [REDACTED] Did we have a rough makeup of our H2 priorities after the SSI team brainstorming sessions?

[REDACTED] (5/22/2023 10:35:36 PDT):

>Yes, I should have the doc out by EOD

[REDACTED] (5/22/2023 13:25:30 PDT):

>Hi @ [REDACTED], what is your estimate of the timeline to update Borderline policy with the changes you are recommending? The context is that I want to know if @ [REDACTED] will have the new policy operationalized in time to build a new version of the SSI borderline classifier with it.

>@ [REDACTED] - what would be the drop dead date for you in H2 to have policy changes finalized so that Ops can train reviewers > you can train the classifier and we can update interventions in H2?

[REDACTED] (5/22/2023 14:15:48 PDT):

>Hi @ [REDACTED] - changes a) which were Cat 6 + 1 secondary indicator (instead of 2) is done. We just need to launch that. @ [REDACTED] talked to you separately about that. Perhaps 2 weeks from now?

>
>Changes b) which were b.1. merging some borderline categories (as per my suggestions in the analysis document) and b.2. which is re-formatting the policy (as today is difficult to read) are [REDACTED] team's jurisdiction now, so I can't provide an estimate myself (though I'll be supporting them / reviewing / counseling) — @ [REDACTED] perhaps you want to weight into this?

[REDACTED] (5/22/2023 14:37:28 PDT):

>@ [REDACTED] My guess is that the changes from a) above itself are enough for a significant revision of the classifier. What do you think?

>I feel we should definitely try to incorporate as many changes as we can. My point is, with the news above, our chances of using an improved signal in interventions in H2 seem quite good.

[REDACTED] (5/22/2023 15:02:35 PDT):

>The classifier yes, though the long pull is generally getting the training data. In this case as we already capture the secondary indicators as part of the current labelling process see

https://www.internalfb.com/intern/icc/tab/decision_trees/tree/edit/272597547114418/

>
>we can retroactively change the decision string (assuming that the label level quality is good enough) and it wouldn't require us to wait for training data. So for a) if we make it by before end of June we should be good.

>
>For b) please share more details about the proposed changes. I want to understand if this would be backward compatible with the existing training data

[REDACTED] (5/22/2023 15:04:43 PDT):

>@ [REDACTED] - can you please point [REDACTED] to your doc with b) changes? Thanks!

[REDACTED] (5/22/2023 15:28:24 PDT):

>@here - we have approval for launching PV demotions on FB feed for non-EU locales. This is contingent to being able to provide review capacity to handle appeal volumes in EU and then launch EU next half. We have to be sure we are able to do this before launch in non-EU this half. @ [REDACTED] - can you please advise on how we can

>1/ estimate appeal volumes

>2/ request for appeal volumes.

>

>cc - @ [REDACTED] @ [REDACTED]

[REDACTED] (5/22/2023 15:32:28 PDT):

>@ [REDACTED] Do we have a POC on the FB side for running experiments?

[REDACTED] (5/22/2023 15:32:56 PDT):

>which experiment? The PV demotion above?

[REDACTED] (5/22/2023 15:43:43 PDT):

>PV demotion and age gating

[REDACTED] (5/22/2023 15:45:19 PDT):

>PV demotions on FB - [REDACTED] (PM) team. Age-gating on FB will not launch this half due to policy dependencies. Lets take this discussion in our chat.

[REDACTED] (5/22/2023 15:58:23 PDT):

>@here - We have a *WIP* H2 roadmap doc. Some points to note:

>1/ I have added items based on the discussion during road-mapping week. I will need help from @ [REDACTED] to add eng centric items discussed in the eng brainstorming. @ [REDACTED] - I will need your help to add ops items discussed like efficiency ones. I feel we should ask ourselves if Precision should remain a theme - I think now its more like keeping the lights running; let me know if you disagree.

>2/ There have been no specific requirements from *regulation* yet. That team is trying their best and we may have some requirements in the coming weeks. We just have to adapt when and if that happens. We have to identify what we will drop in that case.

>3/ Some items will get confirmed after we conclude the High CoV discussion with leads and CIPC review in June.

>

> https://docs.google.com/document/d/1IW5wTtq99MMPrI9_Q6-q920w2RXD2qq_pQKzwulNoK0/edit#heading=h.mo094zb051m9#shareadoc edit

[REDACTED] (5/22/2023 16:01:15 PDT):

>I have yet to work in items based on @Josh's analysis published earlier today; working on it tonight. I recommend taking the time to read it. Lots of insights - <https://fb.workplace.com/groups/1288930768602423/permalink/1483032102525621/>

Josh Latshaw (5/22/2023 16:05:03 PDT):

>Thank you! Apologies about posting in a different closed group. I will be incorporating feedback and will share a note to make it easier in the coming ~2 weeks. In the meantime, Prathish and I are usually pretty quick to respond to join requests, and I'll approve as quickly as I can

[REDACTED] (5/22/2023 17:40:18 PDT):

>Yes! @ [REDACTED] see options 3 & 4 -

https://docs.google.com/document/d/1FUQ_Do8nLRRyWvWvGpc8okhNuqYAf1o_duYJu4ZUVDg/edit