

BEKAH, BOGDAN, ERO, CARMELA, SEDA

THE POTS WORKBOOK

CURIOUS ACADEMICS AT THE IFF

Copyright © 2019 Bekah, Bogdan, Ero, Carmela, Seda

PUBLISHED BY CURIOUS ACADEMICS AT THE IFF

TUFTE-LATEX.GOOGLECODE.COM

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0>. Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

First printing, April 2019



AUTONOMOUS TRAP 001

JAMES BRIDLE

*Dedicated to those Pokémon Go players
that inspired us.*

Introduction to the Workbook

This workbook is made to accompany you during and after the workshop “A guide to revolutionary counter-optimization” at the Internet Freedom Festival, April 2019. Taking as its starting point a framing of optimization and its miscontents, this workshop takes a quick dive into the technologies and applications that exemplify the optimization logic. Participants are invited to examine existing systems and our experiences with them to gauge and broaden our conception and understanding of optimization systems and their harms. We will then explore techniques to expand our political agency under these circumstances. Through these counter-political strategies, participants situate our relationship with these systems and explore the means at our disposal—be they subversion, queering, resistance, or militancy—for engagement within and outside of the optimization regime.

But, that’s a lot of words in one go, and definitely not very optimized, so let’s start!¹

¹ The link to the PDF file of this workbook can be found on our IFF pad: <https://pad.internetfreedomfestival.org/p/1028>

Optimization Systems

Optimization is fundamental to the logic of surveillance capitalism. To elicit optimal socio-economic performance from all connected agents, their actions are perpetually monitored in environments sensed and co-constituted by digital services. The ever expanding cloud-based service ecosystem incorporates populations and environments, be it by capturing personal fitness and productivity at work, the medical analysis of bodies, the rendering and production of objects, or the circulation of material goods. Their predominant ethos tends towards radical transparency and centralized operations, a form of operations control that focuses on financially relevant outcomes rather than the consequences of the processes it invokes.²

Optimization systems externalize risk-by-design as they scale up and down between various “disruptive” services. This strategy, proponents of which like to call “creative disruption”, expands these systems into spaces (yet) void of regulation, accountability, responsibility and due process or simply exploits loopholes around existing regulations, which were the results of long years of social struggle. In this way, they rush to become “new infrastructure” before social, economic, political or environmental damages surface. These “service providers” are averse to journalistic inquiry and democratic oversight, satisfying themselves through reluctant submission to regulation already watered down through relentless lobbying. Optimization here is the *modus operandi* for a global infrastructure that aspires to replace the conventional role of governments with an amalgamation of private for-profit services run through massive interconnected datasets and automated and autonomous processes.

It is by now clear that optimization systems create substantial ‘negative externalities’. Structurally they disregard non-users, non-humans and environments that do not promise to present any potential socio-economic profit. It means that they engender situations where the actions of a targeted group of agents, e.g., consumption, production, and investment decisions, will have “significant repercussions on agents outside of the group”. To ensure adaptivity, these systems rely on constant exploration of alternatives, where also the risks of exploration are externalized to users and their environments. They are set up to sow the benefits of experimentation while users are responsabilised to deal with unknown environmental consequences. Similar to the functionings of globalised finance capital, extracted value is globalized through financialization, negative externalities remain local and have to be dealt with using local resources. The functioning of these systems rely on mass data collection. In the pursuit of more effective inferences about populations and environ-

² Fenwick McKelvey defines optimization as: (1) techniques of logistics and control, (2) discourses legitimating a mathematical state as a solution to social contention

ments, resources and power concentrate in the hands of a few data holders with increasing risks for privacy and institutionalized abuse of power.

Popular discussions of these systems focus on preventing and repairing social risks and harms such as social sorting, mass manipulation, majority dominance, and minority erasure, but often fail to address the possibility of inherent unfairness in the specific understandings of 'optimal'. They also risk to limit their response to consequences for agents included in the model, while agents external to them often bear the greatest risks. Developing a counterpolitics to these systems means addressing the logic embedded in surveillance capitalism. It means trying to come to terms with the interrelated scales and space-time entities affected by this logic.

Enter POTs

PROTECTIVE OPTIMIZATION TECHNOLOGIES are solutions that enable optimization subjects to defend themselves, individually and collectively, from the unwanted consequences of optimization systems. POTs aim to enable those affected by optimization systems to influence, alter, and contest these systems from the outside, rather than depending on the service providers and otherwise technocratic approaches. POTs is a method and a call to action. It is inspired by strategies developed in the wild against real optimization systems, that, for example, strategically feed extra information to the system in order to change its behavior. Smart POTsy initiatives include neighborhood dwellers negatively affected by Waze's traffic redirection that have fought back by reporting road closures and heavy traffic on their streets—to have Waze redirect users out of their neighborhoods; Pokémon users in rural areas that spoof their locations to urban (Pokémon-dense) areas; other users report to OpenStreetMaps—used by Pokémon Go—false footpaths, swimming pools and parks, resulting in higher rates of Pokémon spawn in their vicinity; finally Uber drivers have colluded to temporarily increase their revenue by simultaneously turning off their apps, inducing a local price surge, and turning the app back on to take advantage of the increased pricing.³

Protective optimization technologies (POTs) aspires to put together knowledge from computer science, law, arts and our ingenuity to increase the robustness of counter-optimization initiatives that benefit the "sub-optimal" among us. They respond to optimization systems' effects on a group of people or local environment by finding ways to reconfigure the workings of these systems from the outside. POTs analyze how capture of events (or lack thereof)

³ "Silicon Valley, however, does not exist in an intellectual vacuum; it depends on a certain type of computer science discipline. Therefore, a people's remake of the Valley will require a people's computer science. Can we envision this?" Jimmy Wu

affect users and environments, then manipulate these events to influence system outcomes, e.g., altering optimization models and constraints or poisoning system inputs. To design a POT, we first need to understand the optimization system. What are its user and environmental inputs (U, E) ? How do they affect the capture of events? Which outcomes $O = F(U, E)$ are undesirable for subpopulations or environments? With a characterization of the system, as given by $F(U, E)$, we identify those who benefit from the system and those placed at a disadvantage by defining a benefit function, $B(X, E') : (x, e', Value) \rightarrow value$ that includes both users and non users ($U \subset X$) and affected environments ($E \subseteq E'$). The disadvantaged are those people and environments that reside in local minima of B and are gravely impacted by the system. We then set an alternative output $B(X, E', Value') : (x, e) \rightarrow value'$ the POT aims to achieve. While of course, anyone may be affected negatively by optimization systems, POTs especially aim to provide agency to those who have been historically marginalized, neglected, made vulnerable from forms of administration, extractivism and global logistics.

A POT's benefit function may attend to different goals. It may attempt to "correct" imbalances optimization systems create, i.e., by improving systems' outcome for populations put at an—often historically continuous—disadvantage. Conversely, it may also strategically attempt to reverse system outcomes as a form of protest, highlighting the inequalities these systems engender. This further hints at the subversive potential of POTs. POT designers may concoct a strategy to produce an alternative to benefit function B to contest the authority of optimization systems, challenging the underlying objective functions these systems optimize to and their very *raison d'être*. To do that, a POT may attempt to sabotage or boycott the system, either for everyone or for an impactful minority that are more likely to effect change, leveraging the power asymmetries the POT precisely intends to erode.

An Analytics for Optimization Systems and POTs

The following is a simple framework for capturing the externalities of optimization systems⁴. Please feel free to add questions that you think we should be asking:

We are not the first ones to identify externalities of optimization systems. One could argue that much of the work on algorithmic fairness, accountability, and transparency are mechanisms to identify and address the externalities of optimization systems. However, the proposals made in these frameworks focus on how to get the service provider to achieve fairness, accountability while being trans-

⁴ The analytics are inspired by Queer Analytics proposed by Helen Pritchard, Jara Rocha and Femke Snelting

 Questions for Analyzing Optimization Systems

1. What are these systems optimizing for? When and where is the optimization taking place?
 2. Who benefits?
 3. Who loses or burdens the costs?
 4. What externalities may these systems have, what damage may they cause?
 5. What can we/users/developers/institutions do about these externalities?
-

parent towards the users. In doing so, they assume that the service providers of PETs are trustworthy and are willing to internalize the costs and damage that optimization systems cause. However, news about service providers failing to provide even basic security and privacy, plus the objective to serve profit suggests that relying on service providers is not sufficient. Moreover, service providers may not be able to recognize or want to internalize the costs of harms due to the interactions between multiple optimization systems, e.g., the negative externalities of Uber, Waze and Lyft on surface roads. So, instead, inspired by people fighting back Waze, Airbnb, Uber, Waze etc. we propose that it should also be possible to push back and resist these costs from “outside of the system”. Such systems are called Protective Optimization Technologies.

Basic design principles for POTs

 Design Steps for POTs

1. Name the externality
 2. Identify who or which environments are harmed
 3. Dissect the roots of its cause: what is the technical problem?
 4. Identify the Goal of the POT and the benefactors of the POTs
 5. Create the Solution
 6. Evaluate impact on benefactors and other actors
-

Further Information

The POTs project

More on what we wrote about optimization systems, POTs and our code:

- Blog Post: Questioning the Assumptions behind Fairness Solutions
<https://www.esat.kuleuven.be/cosic/questioning-the-assumptions-behind-fairness-solutions/>
- Academic paper (pre-print): Protective Optimization Technologies
<https://arxiv.org/abs/1806.02711>
- Academic paper (peer-reviewed): Questioning the Assumptions Behind Fairness Solutions
<https://arxiv.org/abs/1811.11293>
- Code: <https://github.com/spring-epfl/pots>
- Slides:
<https://bogdankulynych.me/iff2019/slides.pdf>
- This workbook:
<https://bogdankulynych.me/iff2019/workbook.pdf>

You can also always contact us with your optimization miscontents or POTs ideas!

- Rebekah Overdorf: rebekah.overdorf@epfl.ch
- Bogdan Kulynych: bogdan.kulynych@epfl.ch
- Ero Balsa: ebalsa@esat.kuleuven.be
- Carmela Troncoso: carmela.troncoso@epfl.ch
- Seda Güses: seda@kuleuven.be

Further readings

TECHNOLOGY

The Perfect Selfishness of Mapping Apps

Apps like Waze, Google Maps, and Apple Maps may make traffic conditions worse in some areas, new research suggests.

ALEXIS C. MADRIGAL MAR 15, 2018



A traffic jam in Los Angeles, like always (REUTERS/BRET HARTMAN)

commune

CURRENT ISSUE

ABOUT

CONTACT

Optimize What?

JIMMY WU



Silicon Valley is full of the stupidest geniuses you'll ever meet. The problem begins in the classrooms where computer science is taught.

Figure 1: On the miscontents of traffic optimization apps by Alexis C. Madrigal. <https://www.theatlantic.com/technology/archive/2018/03/mapping-apps-and-the-price-of-anarchy/555551/>

Figure 2: A proposition to put social progress over technological progress defined through optimization by Jimmy Wu. <https://communemag.com/optimize-what/>

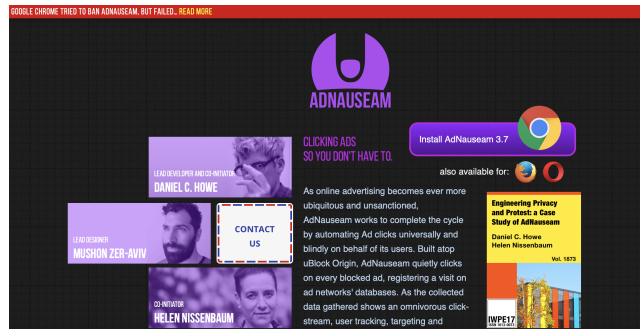


Figure 3: An add on that automates ad clicks universally and blindly on behalf of users intended to make user targeting and surveillance futile <https://adnauseam.io>



New Jersey Town Restricts Streets From Commuters To Stop Waze Traffic Nightmare

May 8, 2018 · 4:01 PM ET

SAMANTHA RAPHELSON



A street sign tells out-of-towners to keep off a side street in Leonia, N.J., where officials are trying to redirect commuters on their way to the nearby George Washington Bridge into New York.
David Porter/AP

Figure 4: A short article on how municipalities resist the externalities of Waze by Samantha Raphelson. <https://www.npr.org/2018/05/08/609437180/new-jersey-town-restricts-streets-from-commuters-t>

Stirring the POTs: Protective Optimization Technologies¹

Seda Gürses, Rebekah Overdorf, Ero Balsa

In the 90s, software engineering shifted from packaged software and PCs to services and clouds, enabling distributed architectures that incorporate real-time feedback from users (Gürses & van Hoboken, 2018). In the process, digital systems became layers of technologies metricized under the authority of optimization functions. These functions drive the selection of software features, service integration, cloud usage, user interaction and growth, customer service, and environmental capture, among others. Whereas information systems focused on storage, processing and transport of information, and organizing knowledge – with associated risks of surveillance – contemporary systems leverage the knowledge they gather to not only understand the world, but also to optimize it, seeking maximum extraction of economic value through the capture and manipulation of people's activities and environments.

The Optimization Problem

The ability of these optimization systems to treat the world not as a static place to be known, but as one to sense and co-create, poses social risks and harms such as social sorting, mass manipulation, asymmetrical concentration of resources, majority dominance and minority erasure.

In mathematical vocabulary, optimization is about finding the best values for an 'objective function'. The externalities of optimization occur due to the way that these objective functions are specified (Amodei et al. 2016). These externalities include:

- 1) Aspiring for asocial behavior or negative environmental ordering (Madrigal 2018, Cabannes et al. 2018),
- 2) Having adverse side effects (Lopez 2018),
- 3) Being built to only benefit a subset of users (Lopez 2018),
- 4) Pushing risks associated with environmental unknowns and exploration onto users and their surroundings (Bird et al. 2016),²
- 5) Being vulnerable to distributional shift, wherein a system that is built on data from a particular area is deployed in another environment that it is not optimized for (Angwin et al. 2016),
- 6) Spawning systems that exploit states that can lead to fulfillment of the objective function short of fulfilling the intended effect (Harris 2018),
- 7) Distributing errors unfairly (Hardt 2014), and
- 8) Incentivizing mass data collection.

Common to information and optimization systems is their concentration of both data and processing resources, network effects, and ability to scale services that externalize risks to populations and environments. Consequently, today a handful of companies are able to amass enormous power.

In the rest of this provocation we focus on location based services (LBS). LBS have moved beyond tracking and profiling individuals for generating spatial intelligence to leveraging this information to manipulate users' behavior and create "ideal" geographies that optimize space and time to customers' or investors' interests (Phillips et al. 2003). Population experiments drive iterative designs that ensure sufficient gain for a percentage of users while minimizing costs and maximizing profits.

For example, LBS like Waze provide optimal driving routes that promote individual gain at the cost of generating more congestion. Waze often redirects users off major highways through suburban neighborhoods that cannot sustain heavy traffic. While useful for drivers, neighborhoods are made busier, noisier and less safe, and towns need to fix and police roads more often. Even when users benefit, non-users may bear the ill effects of optimization.

Users within a system may also be at a disadvantage. Pokémon Go users in urban areas see more Pokémon, Pokéstops, and gyms than users in rural areas. Uber manipulates prices, constituting geographies around supply and demand that both drivers and riders are unable to control while being negatively impacted by price falls and surges, respectively. Studies report that Uber drivers (who work on commission) make less than minimum wage in many jurisdictions.

Disadvantaged users have developed techniques to tame optimization in their favor, e.g., by strategically feeding extra information to the system in order to change its behavior. Neighborhood dwellers negatively affected by Waze's traffic redirection have fought back by reporting road closures and heavy traffic on their streets –to have Waze redirect users out of their neighborhoods. Some Pokémon users in rural areas spoof their locations to urban areas. Other users report to OpenStreetMaps – used by Pokémon Go– false footpaths, swimming pools and parks, resulting in higher rates of Pokémon spawn in their vicinity. Uber drivers have colluded to temporarily increase their revenue by simultaneously turning off their apps, inducing a local price surge, and turning the app back on to take advantage of the increased pricing.

While the effectiveness of these techniques is unclear, they inspire the type of responses that a more formal approach may provide. In fact, these responses essentially constitute adversarial machine learning, seeking to bias system responses in favor of the "adversary". The idea of turning adversarial machine learning around for the benefit of the user is already prevalent in Privacy Enhancing Technologies (PETs) literature, e.g., McDonald 2012. It is in the spirit of PETs that we attend to the optimization problem, i.e., we explore ideas for technologies that enable people to recognize and respond to the negative effects of optimization systems.

Introducing POTs

Protective optimization technologies (POTs) respond to optimization systems' effects on a group of people or local environment by reconfiguring these systems from the outside. POTs analyze how capture of events (or lack thereof) affect users and environments, then manipulate these events to influence system outcomes, e.g., altering optimization models and constraints or poisoning system inputs.

To design a POT, we first need to understand the optimization system. What are its user and environmental inputs (U,E)? How do they affect the capture of events? Which outcomes $O = F(U,E)$ are undesirable for subpopulations or environments? With a characterization of the system, as given by $F(U,E)$, we identify those who benefit from the system and those placed at a disadvantage by defining a benefit function, $B(X, E'): (x,e', \text{Value}) \rightarrow \text{value}$ that includes both users and non users ($U \subset X$) and affected environments ($E \subseteq E'$). The disadvantaged are those people and environments that reside in local minima of B and are gravely impacted by the system. We then set an alternative output $B(X, E', \text{Value}'): (x,e') \rightarrow \text{value}'$ the POT aims to achieve.

A POT's benefit function may attend to different goals (Figure 1). It may attempt to "correct" imbalances optimization systems create, i.e., by improving systems' outcome for populations put at an –often historically continuous– disadvantage. Conversely, it may also strategically attempt to reverse system outcomes as a form of protest, highlighting the inequalities these systems engender. This further hints at the subversive potential of POTs. POT designers may concoct a strategy to produce an alternative to B to contest the authority of optimization systems, challenging the underlying objective functions these systems optimize to and their very *raison d'être*. To do that, a POT may attempt to sabotage or boycott the system, either for everyone or for an impactful minority that are more likely to effect change, leveraging the power asymmetries the POT precisely intends to erode.




Figure 1 Benefit function (top left). POT strategies include redistribution (top right), protest (bottom left), sabotage (bottom right).

Once we select a strategy, we must choose the techniques that implement it. These techniques involve changes to the inputs that users have control over and alterations to

constraints over the objective function to reconfigure event capture (i.e., the system's mechanism of detection, prediction, and response to events). Lastly, we deploy and assess the impact of the POT both in terms of local and global effects on users and environments and tweak it as necessary.

We note that POTs may elicit a counter response from the optimization systems they target, with service providers either neutralizing their effect or expelling POT users. Anticipating these responses may require POT designers to aim for stealth or undetectability, e.g., by identifying minimum alterations to inputs or optimizing constraints to prevent detection.

Discussion

POTs come with moral dilemmas. Some of these compare to concerns raised by obfuscation-based PETs, although these focus on protecting privacy and not protecting populations and environments from optimization. In their work on obfuscation, Brunton and Nissenbaum (2015) highlight four ethical issues: dishonesty, polluted databases, wasted resources and free riding.

Since optimization systems are not about knowledge, we may argue using POTs cannot be judged as dishonesty but as introducing feedback into the cybernetic loop to get optimization systems to recognize and respond to their externalities. POTs are likely to come at a greater cost to service providers and give rise to negative externalities that impact different subpopulations and environments. In fact, all of the harmful effects of optimization systems may be replicated: POTs may have asocial objective functions, negative side effects, etc. One may argue that if optimization is the problem, then more optimization may even come to exacerbate it. Moreover, POTs users may be seen as free riders. These are serious concerns, especially since whichever benefit function B we choose, there will be users who do not agree with or are harmed by the POT. Yet, this problem is inherent to optimization systems' externalities, especially when users are free-riding on non-users or on existing infrastructure.

Banging on POTs: a digital caccerolazo

Optimization history is also one of counter-optimization as evident in the case of search engine optimization or spammers. As optimization systems spread, POTs ensure that counter-optimization is not only available to a privileged few. One could insist that we should work within the system to design better optimization systems. Given service providers' track record in not responding to or recognizing their externalities, POTs aim to explore and provide technical avenues for people to intervene from outside these systems. In fact, POTs may often be the only way users and non-users can protect themselves and secure better outcomes. While short of a revolution, POTs bring people into the negotiations of how their environments are

organized. They also help to provoke a popular response to optimization systems and their many impacts on society.

References:

- Amodèi, Dario, Olah, Chris, Steinhardt, Jacob, Christiano, Paul, Schulman, John, and Mane', Dan. 2016. "Concrete problems in AI safety", *arXiv:1606.06565*.
- Angwin, Julia, Larson, Jeff, Mattu, Surya, and Kirchner, Lauren. 2016 "Machine bias", *ProPublica*, May 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Bird, Sarah, Barocas, Solon, Crawford, Kate, Diaz, Fernando, and Wallach, Hanna. 2016. "Exploring or exploiting? Social and ethical implications of autonomous experimentation in AI", *SSRN Paper ID 2846909*.
- Brunton, Finn and Nissenbaum, Helen. 2015. *Obfuscation: A user's guide for privacy and protest*. MIT Press: Cambridge.
- Cabannes, Théophile, Marco Antonio Sangiovanni Vincentelli, Alexander Sundt, Hippolyte Signargout, Emily Porter, Vincent Fighiera, Juliette Ugirumura, and Alexandre M. Bayen. 2018. "The impact of GPS-enabled shortest path routing on mobility: a game theoretic approach". Presented at Transportation Research Board 97th Annual Meeting. Washington DC, USA. 7–11 January 2018.
- Gürses, Seda and van Hoboken, Joris. Privacy after the Agile Turn. 2018. In Selinger et al. (eds), *Cambridge Handbook of Consumer Privacy*, Cambridge University Press.
- Hardt, Moritz. "How big data is unfair? Medium, September 2014. URL <https://medium.com/@mrtz/how-big-data-is-unfair-9aa544d739de>
- Harris, Malcolm. Glitch capitalism, *New York Magazine*, April 2018. URL <https://nymag.com/selectall/2018/04/malcolm-harris-on-glitch-capitalism-and-ai-logic.html>.
- Lopez, Steve. On one of L.A.'s steepest streets, an app-driven frenzy of spinouts, confusion and crashes. *Los Angeles Times*, April 2018. URL <https://www.latimes.com/local/california/la-me-lopez-echo-park-traffic-20180404-story.html>.
- Madrigal, Alexis C. 2018. "The perfect selfishness of mapping apps." *The Atlantic*, March 2018. <https://www.theatlantic.com/technology/archive/2018/03/mapping-apps-and-the-price-of-anarchy/555551/>.
- McDonald, Andrew WE, Afroz, Sadia, Caliskan, Aylin, Stolerman, Ariel, and Greenstadt, Rachel. 2012. Use fewer instances of the letter "i": Toward writing style anonymization. In *PETS Symposium*, pp. 299–318. Springer.
- Phillips, David, Curry, Michael. 2002. Privacy and the phenetic urge. In D. Lyon (ed.): *Surveillance as Social Sorting: Privacy, Risk, and Digital Discrimination*. NY: Routledge.

¹ We are indebted to Martha Poon for her original framing of the optimization problem and to Jillian Stone for her empirical insights into Pokémon Go. This work was supported in part by the Research Council KU Leuven: C16/15/058; the European Commission through KU Leuven BOF OT/13/070 and H2020-DS-2014-653497 PANORAMIX; and, generously supported by a Research Foundation - Flanders (FWO) Fellowship.

² We disagree with this paper's premise that optimization systems will lead to 'optimal' outcomes, with experimentation as its only potential externality – we appreciate their highlight of the latter.

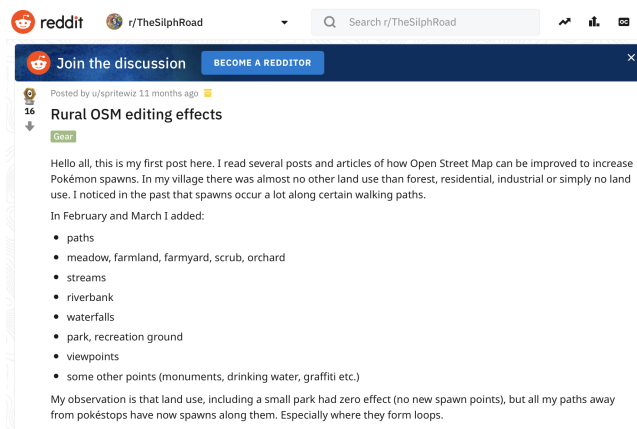


Figure 5: A reddit thread on how to get more Pokémon in “sub-optimal” locations like poor neighborhoods and rural areas by manipulating Open Street Maps <https://tinyurl.com/y5sn6tpt>

