On the Practicalities of Robots in Public Spaces

Cindy M. Grimm and Kristen Thomasen*

*UPDATED DRAFT. *A note to our readers – this draft is current as of September 21, 2021. Please note that if you looked at our earlier draft from August, there have been some substantial revisions since then.*

Introduction

In this paper, we examine laws governing sidewalk robots with a view to understanding what lawmakers are doing - intentionally and (likely) unintentionally - when they adopt these laws. We look at what kind of rules are actually being created for sidewalk robot use and consider, through both legal and technical lenses, what impact the rules will really have in practice on the public urban environments where these technologies are operating. In doing all this, we are also thinking about how legal, civic, and technical experts can work across their areas of expertise to better understand what kinds of rules are actually needed and what some of the impacts of different rules might be, in an effort to understand how lawmakers might design more thoughtful (and hopefully, effective) laws. We encountered questions such as: when is it better to have a pre-emptive law to address an issue raised by a new technology and when is it better to have reactive law-making through the courts; how can lawmakers and technologists understand the costs and tradeoffs in regulating technology (or, can they even understand this in advance); if costs and tradeoffs are understood, how should lawmakers balance and prioritize different and possibly competing interests; and if lawmakers should pre-emptively consider potential re-purposing of deployed robotic technology? We encountered more questions than we could answer when considering the potential impacts of regulating a robot system that relies on public infrastructure, and interacts with people and the built environment. Through all this, though, we were able to identify some lines of questions that need to be addressed explicitly in the context of sidewalk delivery robot regulation, but that so far seem to be overlooked.

The paper starts with an overview of the laws in the U.S. that purport to regulate sidewalk robots. We then consider what kinds of technical requirements are created through different legal objectives, and through this, we identify a number of high stakes tradeoffs in interests that raise questions about how to deal with sidewalk robots through a legal lens, and even, whether it's worth adopting this technology given the costs it might bear, and what benefits the public actually gets out of it. For example, a requirement to make the technology as safe as possible also requires significant data collection; a requirement to keep sidewalk robots out of the way of the public on sidewalks might require significant infrastructural

*Extending enormous gratitude to Allard Law student Uma Wu for her phenomenal assistance on the legal history section of this paper.

investments to widen sidewalks or create separate lanes. We outline who the relevent stakeholders are, and what they should bring to the conversation. We conclude the paper with some reflections on what all of this tells us about law-making and novel technologies that will use public shared spaces, and how to think about the inherent tradeoffs that will always be involved.

In this paper, we're specifically concerned with what the actual functionality of a sidewalk robot means for law-makers and for the people who will encounter these robots on city sidewalks. Law-makers should also be concerned more broadly with what it means for the public, and the public interest, for these technologies to be entering into a shared space, like a public sidewalk. That broader discussion has been touched on by previous We Robot papers, and is beyond the scope of our focus here.[1] We also limit the conversation to non-malicious use-cases. Although "hacked" robots and deliberate misuse of robots in public spaces for surveillance/stalking is a valid problem, it is also outside of the scope of this paper.

For our discussion, we will define a "sidewalk robot" to be a robot that is designed to maneuver along sidewalks and/or bike lanes, sharing this public space with humans. Typical tasks for such a robot include commercial deliveries, but could also include monitoring, autonomous luggage, trash collection, or last-mile delivery.

A Brief Legal History of Sidewalk Delivery Robot Regulation

Sidewalk delivery robots have been a subject of legislation in the US since 2017. As of 2020, 12 US states had some state-wide sidewalk robot legislation, and at least three other states had introduced and debated proposals to legislate.[2] In 2021 (at the time of writing), 14

---

[1] See for example, the Urban Robotics panel on Saturday April 14, 2018 at the Stanford We Robot (with Jesse Woo, Jan Whittington, Ronald Arkin, and Kristen Thomasen: https://conferences.law.stanford.edu/werobot/agenda/) and Mason Mark's paper on "Robots in Space" at the 2019 We Robot in Miami (https://robots.law.miami.edu/2019/wp-content/uploads/2019/04/Marks_Robots-in-Space.pdf)

[2] Y, Gerald. "U.S. States That Have Legalized Personal Delivery Devices (PDDs) or Last-mile Autonomous Delivery Robots", (11 December 2020), online: *INSIGHTS by RList* <https://insights.rlist.io/2020/12/states-that-legalized-pdds.html>. The article says 14, but there are actually only 12.

A bill introduced in Missouri died in committee. The Colorado bill is postponed indefinitely as of 2021. The Bill in Michigan passed in the Senate (22-16) and was under consideration by the House as of September 10, 2020 (https://www.michiganvotes.org/2020-SB-892).

additional sidewalk delivery robot bills have been introduced in states across the U.S., resulting in six new states with state-wide laws, at least one update to an existing state law, and seven other state bills failing to pass.[3] There are also a range of municipal-level ordinances addressing sidewalk robots.[4] In addition to legislation, some private entities have internal rules for the operation of sidewalk robots on their property. These entities are predominantly found in the US and function within a similar framework to the locally applicable legislation.

The focus of sidewalk robot regulation - in both legislation and in rules set by private entities – has been on several aspects of their design and operation, including size, maximum speed, and purpose of operation,[5] as well as whether the sidewalk robot must be monitored or controlled by a human operator.[6] These parameters are largely consistent across jurisdictions, with some notable exceptions. A number of seemingly important parameters are also currently absent from the regulations. We highlight some of those in this section. First, we briefly review which jurisdictions have regulated, and then we consider how states have approached regulating – what components of a sidewalk robot system have caught the attention of law-makers, and what has been overlooked?

As of 2021, no Canadian jurisdiction has regulated sidewalk robots. There are a variety of potential reasons for this disparity, including wider-scope problems such as regulatory zeal potentially rendering Canadian businesses uncompetitive, to narrower-scope causes for concern such as past incidents of "mass noncompliance" when it comes to new technology in that vein. See e.g. Powell, Naomi. "Remember the internet of the '90s? That's what Canada's outdated data protection laws were meant to handle", (12 December 2018), online: *Financial Post* <https://financialpost.com/technology/remember-the-internet-of-the-90s-thats-what-our-outdated-privacy-rules-were-built-to-handle>; "Shared e-scooters to be banned in Montreal in 2020 | CBC News", (19 February 2020), online: *CBC News* <https://www.cbc.ca/news/canada/montreal/scooters-banned-1.5468206>

[3] This site provides updates on state law regarding sidewalk delivery robots: https://www.pedbikeinfo.org/resources/resources_details.cfm?id=5314. See also, Indiana (came into effect April 1, 2021, House in favour 95-0; Senate in favour 36-12-2): http://iga.in.gov/legislative/2021/bills/house/1072?__cf_chl_jschl_tk__=pmd_wdc_UIkUFjmJ8hgLnCfo1mzKWxC4t0fsYCAnI0gaSlI-1631911293-0-gqNtZGzNAjujcnBszQNR#digest-heading and https://legiscan.com/IN/votes/HB1072/2021; Louisiana law came into effect in June 2021 (House in favour 101-0; Senate in favour 34-1): https://legiscan.com/LA/text/SB147/2021 Minnesota has a Bill under consideration: https://www.revisor.mn.gov/bills/text.php?number=HF3621&type=bill&version=0&session=ls91&session_year=2020&session_number=0

[4] San Francisco (https://codelibrary.amlegal.com/codes/san_francisco/latest/sf_publicworks/0-0-0-48516); Madison WI (https://library.municode.com/wi/madison/codes/code_of_ordinances?nodeId=COORMAWIVOIICH11--19_CH12VECO_BIPLVE_12.753REPEDEDE); District of Columbia (https://code.dccouncil.us/dc/council/laws/22-137.html)

[5] Va Code Anna § 46.2-100 (2020) [Virginia Law]

[6] US, SB 148, *An Act relating to personal delivery devices and providing a penalty*, 2017-148, Reg Sess, Wis, 2017 (enacted) [Wisconsin Law].

Jurisdictions and Entities with Sidewalk Robot Regulation

Between 2017 and 2021, at least 18 US states have passed legislation seeking to integrate sidewalk robots onto public sidewalks, roads, bike lanes, and/or intersection crossings.[7] The early laws were the product of several years of lobbying from a variety of interest groups that stood to profit from mass deployment of sidewalk robots, including but not limited to Starship robots - the lobbyists behind the state proposals for legislation in Virginia, Idaho, and Wisconsin[8]. Lobbying and involvement from various sidewalk robot companies continues today. For example, FedEX introduced the bill that was considered in Minnesota in 2021.[9]

In 2017, Virginia[10], Idaho[11], Wisconsin[12], Florida[13], and Ohio[14] became the first five states to have legislation pertaining to sidewalk robots, in that order of adoption. The bills were either passed unanimously[15] or with an overwhelming majority[16], with the exception of the bill in Ohio, which encountered some resistance in the House in the form of a 58-37 split along party lines in the final vote, with most Yeas by Republicans and most Nays by Democrats[17]. The Virginia bill passed in a little over four months[18], and the Florida[19] and Ohio[20] bills in a little over three.

---

[7] See e.g. Y, "U.S. States that have Legalized", *supra* note 2 and https://www.pedbikeinfo.org/resources/resources_details.cfm?id=5314

[8] Glaser, April. "Florida is now the fourth state to permit delivery robots on sidewalks", (26 June 2017), online: *Vox* <https://www.vox.com/2017/6/26/15877278/florida-law-unmanned-delivery-robot-sidewalk-starship-technologies>

[9] http://www.dot.state.mn.us/automated/docs/personal-delivery-device-white-paper.pdf; https://www.revisor.mn.gov/bills/text.php?number=HF3621&type=bill&version=0&session=ls91&session_year=2020&session_number=0

[10] Glaser, April. "Virginia is the first state to pass a law allowing robots to deliver straight to your door", (1 March 2017), online: *Vox* <https://www.vox.com/2017/3/1/14782518/virginia-robot-law-first-state-delivery-starship>

[11] Glaser, April. "Idaho is the second state to allow unmanned robots to deliver to your front door", (27 March 2017), online: *Vox* <https://www.vox.com/2017/3/27/15075048/idaho-unmanned-robots-law-delivery-starship>

[12] Glaser, April. "Wisconsin is now the third state to allow delivery robots", (22 June 2017), online: *Vox* <https://www.vox.com/2017/6/21/15850906/wisconsin-law-delivery-robots-starship>

[13] Wisconsin Law, *supra* note 4.

[14] Glaser, April. "Ohio is now the fifth U.S. state to permit delivery robots on sidewalks", (5 July 2017), online: *Vox* <https://www.vox.com/2017/7/5/15916688/ohio-fifth-state-delivery-food-robots-starship-law>

[15] See Wisconsin Law *supra* note 4; US, HB 1027, *An Act relating to unmanned devices*, 2017, Reg Sess, Flo, 2017 (enacted) [Florida Law].

[16] DeSteph, Bill R., online: *Legislative information system* <https://lis.virginia.gov/cgi-bin/legp604.exe?171%2Bsum%2BSB1207>; US, HB 204, *An Act Relating to Personal Delivery Devices*, 2017-64, Reg Sess, Ida, 2017 (enacted) [Idaho Law]; US, HB 566, *An Act Relating to Motor Vehicles*, 2020-65, Reg Sess, Ida, 2020 (enacted).

[17] US, HB 49, *FY 2018-2019 operating budget*, 132st Gen Assem, Reg Sess, Ohio, [2017] (enacted) [Ohio]

[18] See *supra* note 13 (DeSteph).

[19] See Wisconsin Law *supra* note 4.

[20] See *supra* note 14.

In 2018, Utah[21] and Arizona[22] also passed legislation on sidewalk robots. The Utah bill was passed unanimously[23], whereas the Arizona bill encountered slight resistance in the House, passing with 33 Yeas and 24 Nays at the third reading.[24] This vote yet again split along party lines, with predominantly Republican proponents and Democratic detractors.[25] Both bills passed extraordinarily quickly, at little over a month before ratification.[26]

In 2019, Washington (state)[27] and Texas[28] passed legislation on sidewalk robots. The bill passed unanimously in Texas[29] and encountered only slight resistance in the house in Washington state, receiving 94 Yeas and 3 Nays in the House before passing.[30] The Texas bill passed in around 3 months,[31] whereas it took almost 8 months in Washington.[32]

In 2020, North Carolina[33] and Pennsylvania[34] also passed legislation on sidewalk robots. A bill on the subject was also introduced in Missouri, but has since died in committee.[35] A bill was introduced in Colorado in January of 2020, but debate was postponed indefinitely in February that year.[36] The North Carolina bill was passed unanimously except for a single

---

[21] US, HB 227, *Personal Delivery Devices Amendments*, 2020, Gen Sess, Utah, 2020 (enacted) [Utah Law]

[22] US, HB 2422, *An Act Relating to Transportation Devices*, 2018-53, 2nd Reg Sess, Ariz, 2018 (enacted) [Arizona Law].

[23] See Utah Law *supra* note 18.

[24] See Arizona Law *supra* note 19.

[25] *Ibid.*

[26] See Utah Law *supra* note 18. See Arizona Law *supra* note 19.

[27] US, HB 1325, *An Act Relating to the regulation of personal delivery devices*, 2019-66, Reg Sess, Wash, 2019 (enacted) [Washington Law]

[28] US, SB 969, *An Act relating to the operation of personal delivery and mobile carrying devices*, 2019-86, Reg Sess, Tex, 2019 (enacted) [Texas Law]

[29] *Ibid.*

[30] See Washington Law *supra* note 24.

[31] See Texas Law *supra* note 25.

[32] See Washington Law *supra* note 24.

[33] US, SB 739, *An Act to Define and Regulate Personal Delivery Devices*, 2020-73, Reg Sess, NC, [2020] (enacted) [North Carolina Law]

[34] US, SB 1199, *An Act amending Title 75 (Vehicles) of the Pennsylvania Consolidated Statutes, in general provisions, further providing for definitions and providing for personal delivery devices; and making editorial changes*, 2020, Gen Assem, Penn, 2020 (enacted) [Pennsylvania Law]

[35] US, HB 2290, *An Act to Allow personal delivery devices to operate on sidewalks and roadways*, [2017-148], Reg Sess, Mo, [2017] (enacted) [Missouri Bill]

[36] US, SB 20-092, *An Act Concerning the regulation of self-propelled devices used to deliver cargo, and, in connection therewith, specifying standards for the operation of robotic devices within pedestrian areas and on highways*, 2020, Reg Sess, Colo, 2020 (lost) [Colorado Bill]

Nay in the house[37]. There was no data on the votes for the Pennsylvania bill.[38] The North Carolina bill passed in just under two months[39], whereas the Pennsylvania bill passed in just over five.[40]

In 2021, Indiana, Iowa, Oklahoma, Maryland, Arkansas, and Louisiana added to the group of states with state-wide sidewalk delivery robot laws. At least some of these states explicitly cited the pandemic as the catalyst for passing laws to allow autonomous testing and delivery.[41]

There are also some jurisdictions that have been legislatively resistant to the operation of sidewalk robots thus far. A bill on the subject was defeated after its introduction in Kansas,[42] and the city of San Francisco banned the operation of sidewalk delivery robots without a permit in 2017.[43]

Canadian Legislation

As of 2020, there is no federal, provincial, or municipal regulation in Canada on the subject of sidewalk robots, even though there are some sidewalk robots operating in Canada. There is no mention of sidewalk robots or the like in Transport Canada's schedule of proposed regulations,[44] nor is the subject discussed in the strategic plans for Transport Canada up to 2030[45]. On a provincial level, there is no discussion of sidewalk robots in Toronto (where

---

[37] See North Carolina Law *supra* note 30.

[38] See Pennsylvania Law *supra* note 31.

[39] See Texas Law *supra* note 25.

[40] See Pennsylvania Law *supra* note 31.

[41] See e.g. Oklahoma, https://www.natlawreview.com/article/autonomous-bots-personal-delivery-devices-oklahoma

[42] Marshall, Aarian. "Amazon and FEDEX push to Put delivery robots on your sidewalk", (25 August 2020), online: *Wired* https://www.wired.com/story/amazon-fedex-delivery-robots-your-sidewalk/. And see https://airtable.com/shrKZUEh0kKbzZ8fQ/tblzpgo81Mue25I5D?backgroundColor=red&viewControls=on – bills have failed to pass in other states including New Hampshire, Rhode Island, Oregon, Minnesota and Massachusetts.

[43] San Francisco Public Works. "Autonomous delivery devices", online: *Autonomous Delivery Devices | Public Works* <https://www.sfpublicworks.org/services/permits/autonomous-delivery-devices>

[44] Transport Canada. "Schedule of proposed regulations", (20 July 2021), online: *Transport Canada* <https://tc.canada.ca/en/road-transportation/motor-vehicle-safety/schedule-proposed-regulations>

[45] Transport Canada. "Transportation 2030: A strategic plan for the future of transportation in Canada", (25 November 2019), online: *Transport Canada* <https://tc.canada.ca/en/initiatives/transportation-2030-strategic-plan-future-transportation-canada>

Tiny Mile Robots's delivery robot, Geoffrey, is already in operation[46]) and the BC Road Safety Strategy as of 2020[47] (where electric scooters were not allowed until 2019[48] out of concerns of accountability in instances of misuse like those arising in Ottawa following rollout of Bird and Lime scooters[49]).

Despite this barren policy-scape, disparate and small-scale policy discussions on the subject are in fact taking place in Canada. The Centre for Integrated Transportation and Mobility (CITM) based in Ontario has published a white paper advocating for "an international standard to regulate and manage sidewalk robots"[50]. Harmonize Mobility, a partner of CITM, is currently in the process of drafting an international technical standard titled "The global standard for the future sidewalk and curb", which includes discussions on necessary regulations for robotic entities operating on the sidewalk[51].

Private Entities and Sidewalk Delivery Robot Regulation

One of the most ubiquitous applications of sidewalk robots in non-municipal areas is in universities. As of 2021, 16 college campuses in the US are permitting a variety of robots to deliver food and other sundries to their students[52]. Notably, this includes several universities currently in states where there is no other regulation on the deployment of sidewalk robots, including UCLA.[53] The restrictions on the operation of sidewalk robots on these campuses

---

[46] Shea, Courtney | Photography By Daniel Neuhaus | 09/26/2020. "'We liken it to the invention of the washing machine': A Q&A with the creator of Geoffrey, Toronto's adorable new delivery robot", (26 September 2020), online: *Toronto Life* <https://torontolife.com/food/tiny-mile-robotics-geoffrey-robot-delivery-courier/>
Granted, Geoffrey isn't autonomous and still requires a human operator, but it fits within US definitions of PDD.

[47] Ministry of Public Safety and Solicitor General & RoadSafetyBC. "BC Road Safety Strategy", (8 July 2021), online: *Province of British Columbia* <https://www2.gov.bc.ca/gov/content/transportation/driving-and-cycling/roadsafetybc/strategy>

[48] "Electric kick scooters could be legal in 6 B.C. municipalities this summer | CBC News", (23 March 2021), online: *CBC News* <https://www.cbc.ca/news/canada/british-columbia/escooters-scoot-scoot-bc-pilot-1.5960949>

[49] "E-scooters proving popular - but they're not without their critics | CBC News", (5 September 2020), online: *CBC News* <https://www.cbc.ca/news/canada/ottawa/popular-e-scooters-also-causing-problems-1.5713182>

[50] Harmonize Mobility. "Sidewalk and curb", online: *Harmonize Mobility* <https://harmonizemobility.com/sidewalkandcurb/>

[51] *Ibid*

[52] Barack, Lauren. "Robots delivering food to college campuses this fall", (28 January 2021), online: *Gearbrain* <https://www.gearbrain.com/robots-deliver-food-college-campus-2646352550.html>

[53] Danesh, Noah. "UCLA restaurants SERVE campus community with autonomous delivery robots", online: *Daily Bruin* <https://dailybruin.com/2021/02/23/ucla-restaurants-serve-campus-community-with-autonomous-delivery-robots>, see also Repko, Melissa. "On University of Texas at DALLAS' GROWING campus, meal-delivering robots make splashy debut", (25 December 2019), online: *Dallas News* <https://www.dallasnews.com/business/technology/2019/12/26/on-university-of-texas-at-dallas-growing-campus-meal-delivering-robots-make-splashy-debut/>.

are particularly stringent compared to even the most constraining regulations that have been legislated at the state level. For example, the speed limit of Starship Robots operating at the University of Texas at Dallas is 4mph[54], whereas at the state level the speed limit is 10mph on pedestrian areas and 20mph on roadways and highways (notwithstanding municipal restrictions)[55]. Private and university campuses might be a site of important regulation for these devices going forward, though because these are not all public institutions gaining access to their rules and regulations can be difficult. These are spaces worth noting though, as some of our observations below will also be relevant to the less formal law making processes that guide operations in these shared spaces.

Subjects of Regulations

Since there is no Canadian regulation on the subject, the following sections focus on the regulation in the US. The below sections break down what attributes of sidewalk robot design and use have been regulated by state laws. We will later consider what the legal focus on these attributes means in terms of their design and the potential effectiveness of these laws in meeting the (presumed) goals around pedestrian and public safety, flow of sidewalk movement, sidewalk access for pedestrians, infrastructure, *etc.*

Maximum Weight

Every regulation noted above contains some weight limit for these devices. Most of the states have a weight limit for unladen devices of around 100lbs[56], with exceptions in Virginia, North Carolina, and Pennsylvania with weight limits up to 500lbs[57], 500lbs[58], and 550lbs[59] respectively.

For context, the most ubiquitous sidewalk delivery robots in operation in the US include the Starship Robot (50lbs) and the Amazon Scout (100lbs)[60]. The KiwiBot is also in operation on several university campuses[61], but its weight is not available online.

---

[54] *Ibid.*

[55] See Texas Law *supra* note 25.

[56] See Y, "States that have Legalized" *supra* note 2.

[57] See Virginia Law *supra* note 3.

[58] See Texas Law *supra* note 25.

[59] See Pennsylvania Law *supra* note 31.

[60] Omondi, Derick. "Autonomous delivery VEHICLES DIMENSIONS & Drawings", online: *Dimensions & Drawings | Dimensionscom* <https://www.dimensions.com/collection/autonomous-delivery-vehicles>.

[61] See Barack, "Robots Delivering food" *supra* note 48.

Speed Limit

Similar to the weight limit, every regulation examined here has specified a maximum speed for the operation of sidewalk robots. Most of the states set the speed limit at 10mph on the sidewalk[62] and up to 20-25mph on the road[63], with the exception of Washington, which has the lowest speed limit for sidewalk operation at a mere 6mph[64].

For context, the current max speed of the Starship Robot is 3.7mph, the Amazon Scout at 15mph, and KiwiBot at 1.5mph[65]. Typical speed restrictions for electric scooters are also in the 20mph range.

Safety and Accountability Measures

Legislation for sidewalk robots detail a variety of safety and accountability measures meant to ensure that individual devices can be tracked, and held accountable for infractions.[66]

Aside from compliance to local traffic regulations and the aforementioned speed limit, the most common measure is some form of identification for the robot and its operator. For example, the Virginian regulation requires a unique identifying device number and a means of identifying the operator "that is in a position and of such a size to be clearly visible"[67], while the Floridian legislation requires that devices "include a plate or marker that has a unique identifying device number and identifies the name and contact information of the personal delivery device operator"[68]. Tennessee also requires the robot to have some method of displaying its capacity (or lack thereof) for carrying hazardous materials[69].

---

[62] For comparison, the average walking speed of a human is 3-4mph, running is 6-8mp, and bicycling is between 10 and 15mph (up to 22mph for race speeds) on flat ground,

[63] See *supra* note 1.

[64] See Washington Law *supra* note 24.

[65] See Omondi, "Dimensions and Drawings" *supra* note 56.

[66] The Missouri bill says explicitly that the means of identification should be accessible such that the device can be traced if people wanted to.

- "A personal delivery device shall have all of the rights and responsibilities as a pedestrian, must display a unique identifying number, and *be equipped to identify the personal delivery device operator*."

Other states require ID but don't say the purpose like Colorado, Tennessee, Virginia, and Wisconsin.

[67] Va Code Ann § 46.2-908.1:1 (2020).

[68] See Florida Law *supra* note 12.

[69] US, HB 2365, *An Act relative to personal delivery devices*, [2017-148], Reg Sess, Tenn, [2017] (enacted) [Tennesee Law]

There are also a number of mechanical safety requirements for sidewalk robots. Braking mechanisms are explicitly required in some states, including Wisconsin[70] and Utah[71]. In the same vein, lights are required for operations between sunset and sunrise in certain states, including Ohio[72] and Texas[73].

Finally, an accountability measure of note is the possession of an insurance policy by the robot operator. Florida[74] and Texas[75], among other states, require "an insurance policy that includes general liability coverage of not less than $100,000 for damages arising from the operation of the device"[76].

There have been very few publicly reported examples of actual collisions involving sidewalk robots to date. But notably in one instance that received some public attention, identification and insurance proved problematic.

A Starship robot collided with the side of Jisun Mok's car in Texas in September, 2020 causing $2,600 damage to the car. Mok struggled to receive compensation for the damages.

The first challenge Mok faced was that there was no phone number on the robot at that time of the occurrence. According to the Texas legislation in the area (which was effective in 2019, and this incident occurred in 2020), sidewalk robots must "be equipped with a marker that clearly states the name and contact information of the owner and a unique identification number". This requirement was either unfulfilled or was not fulfilled in such a way that the information could be identified.

The second challenge Mok faced was a legal challenge - whether the robot counted as a "vehicle" for the purposes of municipal assistance. Mok sought police assistance, but they were unsure if the robot constituted a vehicle, such that this amounted to a vehicular accident and came within police jurisdiction. Mok also reached out to the city but to no avail. Based on the state legislation, "a personal delivery [device …] is not considered to be a vehicle".

---

[70] See Wisconsin Law *supra* note 4.

[71] See Utah Law *supra* note 18.

[72] See Ohio *supra* note 14.

[73] See Texas Law *supra* note 25.

[74] See Florida Law *supra* note 12.

[75] See Texas Law *supra* note 25.

[76] *Ibid.*

The third problem was proof of fault, for the purposes of an insurance claim. The insurance company told Mok she needed to prove that the accident was not her fault, but Starship refused to release the video taken by the robot during the incident. Starship claimed that the video taken by the robot was its property, so Mok had no right to access it.

Mok was finally able to receive compensation after she went to the local news network, and the network reached out to Starship. There was no record of Mok filing a lawsuit, it appears that Starship just agreed to pay the damages.

All of these barriers arose in spite of legislation that was meant to aid with accountability and mitigate these problems, and in some ways, barriers actually arose *because* of the legislation. This example points to some of the unintended impacts of current sidewalk robot laws, which we break down in more detail in the sections below.

Additionally, Emily Ackerman shared her first-hand account of how sidewalk robots can create unsafe sidewalk conditions for people who use wheelchairs.[77] She details how the presence of a sidewalk robot and its limited capacity to respond to pedestrian congestion at a intersection crossing impeded her ability to safety cross a street. The challenges faced by Mok might even be exacerbated where a pedestrian has not experienced a physical collision (and personal injury or property damage) but has been placed in an unsafe position by the device. Inferring from Mok's and Ackerman's experiences one might foresee situations where, in spite of laws prioritizing pedestrian right of way, pedestrians might face difficulty or lack of safety when attempting to exert their right of way and an associated lack of recourse. Impediments to pedestrian flow and safety could arise for seemingly practical or technical reasons when viewed from a design perspective (e.g. an overly cautious robot), that might even be supported by law. We consider some of these challenges below.

Presence of Human Operators (Permissible Degree of Independent Operation)

There are a range of regulatory requirements for the presence of a human operator for a sidewalk robot operating on the road, ranging from a requirement for full supervision by a human operator, to unsupervised operation. On the full supervision end, Wisconsin sidewalk robot operators "must control or monitor the navigation and operation of the [sidewalk

---

[77] https://www.nbcdfw.com/news/nbc-5-responds/nbc-5-responds-what-happens-if-youre-in-a-crash-with-a-robot/2442345/; https://www.bloomberg.com/news/articles/2019-11-19/why-tech-needs-more-designers-with-disabilities

robot]" at all times while operating on a sidewalk or crosswalk[78]. Closer to the middle of the spectrum, in Colorado's proposal robots could remain within "the remote support and supervision of a human"[79], which implies some limited capacity for the robot to operate without the direct input of a human operator. Similarly, in North Carolina robots only need to be "monitored by an operator *who is able to* exercise remote control over the navigation and operation of the personal delivery device"[80]. This can be distinguished from the Wisconsin legislation, which explicitly requires remote control. On the other end of this spectrum, the Tennessee legislation does not include any explicit requirement for the robot to be operated by a person at all[81]. This is supported by the fact that, in the summary of the bill, the definition of a sidewalk delivery robot includes the stipulation that it is "capable of navigating with or without the active control or monitoring of a natural person", as well as the absence of any requirement for such control or monitoring in the body of the regulation[82].

The presence or absence of human control, as permitted or required by law, raises a range of legal and technical questions including how the presence/absence of a person affects an accountability analysis should something go wrong,[83] to how quickly a company could deal with a robot that breaks down in the middle of a sidewalk,[84] to operational questions about whether it's worth investing in a robot system to deliver items in the first place?

Right of Way, and Pedestrian Rights for Sidewalk Robots

Generally, states have required robots to comply with traffic laws and yield to other traffic, which typically includes pedestrian traffic. In Virginia and Colorado, sidewalk robots are required to yield the right of way and "otherwise not unreasonably interfere with" pedestrians[85].

Several states have granted pedestrian's rights to sidewalk delivery robots, either explicitly or implicitly. In Pennsylvania and Ohio, pedestrian rights are granted implicitly by expanding

---

[78] See Wisconsin Law *supra* note 4.

[79] See Colorado Bill *supra* note 33.

[80] See Texas Law *supra* note 25.

[81] See Tennessee Law *supra* note 65.

[82] *Ibid.*

[83] See e.g. Madeleine Clare Elish, "Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction" (https://estsjournal.org/index.php/ests/article/view/260)

[84] See e.g. Sally Applin, "Delivery Robots Will Rely on Human Kindness and Labor" https://www.vice.com/en/article/ne98x7/delivery-robots-will-rely-on-human-kindness-and-labor

[85] Va Code Ann § 46.2-904 (2020) [Virginia Code]; See Colorado Bill *supra* note 33.

the definition of Pedestrian to encompass sidewalk robots.[86] In Missouri's proposed Bill and Arizona's law, sidewalk delivery robots are explicitly given the "rights and duties applicable to a pedestrian under the same circumstances" [87] with certain additional requirements, such as being required to yield to other pedestrians or being equipped with lighting and identification.[88] The attachment of pedestrian rights to autonomous devices like sidewalk robots raises the important question of what recourse sidewalk robot owners and pilots may have against infringement of the device's "pedestrian" rights. The outcome of these legislative decisions remains to be seen.

Restrictions by Local Ordinances

Certain states have restricted the imposition of local ordinances to manage operation of sidewalk robots, and others have empowered local ordinances to override state-wide regulations where appropriate. States like Virginia[89] and North Carolina[90] have restricted regulation by legislative bodies at the county and municipal levels, with Virginia restricting regulation entirely to the state level[91], and North Carolina disallowing local regulations that forbid the operation of sidewalk robots[92]. States like Idaho[93] and Florida[94] have explicitly enabled such bodies to regulate the operation of sidewalk robots to suit local needs. States like Wisconsin[95] and Washington[96] have also been implicitly enabled to regulate sidewalk robots with legislation that says that state-level regulation of sidewalk robots is 'subject to' any applicable local ordinance. Most states allow some degree of local regulation, so long as it does not contradict the state-level legislation.

Overall, the state-wide laws are relatively similar, at least in the scope of attributes of sidewalk robot use and design that they consider. As we mention, there have not yet been many instances of legal conflict involving sidewalk robots that have either made it to court or been the subject of media coverage. So instead of considering legal cases, we continue our analysis through the lens of the technology itself, and consider what these laws really require

---

[86] See Pennsylvania Law, supra note 31 and Ohio Law, Supra note 14.

[87] See Missouri Bill supra note 32 and Arizona Law supra note 19.

[88] Missouri Bill, *ibid.*

[89] *Ibid.*

[90] See Texas Law *supra* note 25.

[91] See Virginia Code *supra* note 80.

[92] See Texas Law *supra* note 25.

[93] See Florida Law *supra* note 12 (HB 566).

[94] See Florida Law *supra* note 12.

[95] See Wisconsin Law *supra* note 4.

[96] See Washington Law *supra* note 24.

for the design and capabilities of the robots, whether the laws can be effective at achieving goals related to public safety and the public interest (however that might be defined), and at what cost (financially, and socially)?

Overview: How to analyze a robot's (potential) capabilities

Having considered how sidewalk robots have been regulated, we now turn to considering how sidewalk robots actually work. Later, we'll bring these two sections together to think about how legal and technical parameters and expertise coincide (or don't…), and what this means for human-robot interactions in public spaces like sidewalks.

There are many definitions of a robot[97] but for the purposes of this paper we'll think of a robot as having three components: 1) Sensors, which measure physical properties of the world, 2) An internal representation of the world with human-understandable semantic labels (eg, car, crosswalk) in the form of numbers and symbols, and 3) A physical body with actuators that can move the robot. Algorithms (code) plus data enable 1) Processing the raw sensor data into the internal representation and 2) Deciding how to move the actuators based on the internal representation and the goal or task.

---

[97] Ask 10 roboticists what a robot is and you'll get 15 answers. At least.

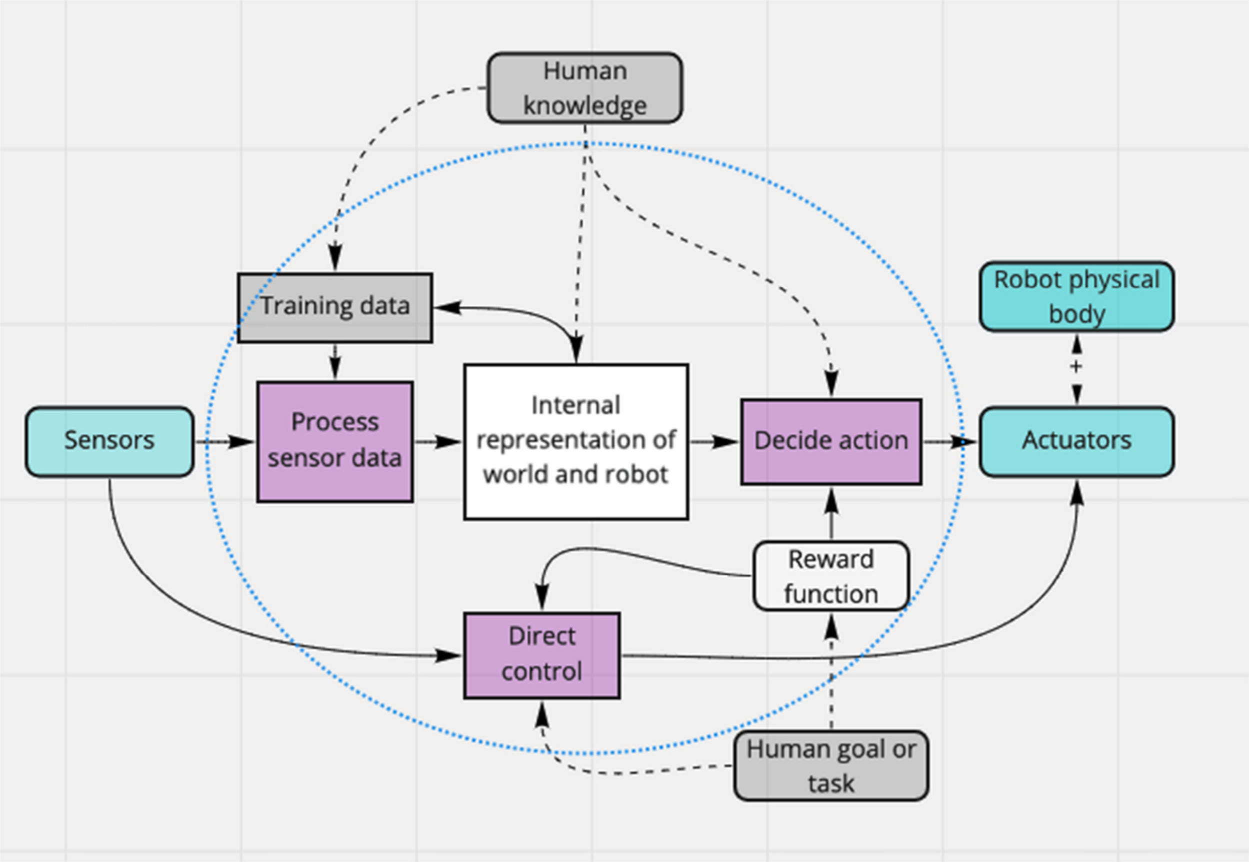

Some observations

- Everything inside the dotted blue line is numbers, symbols, and code.
- Sensors do not have to be inside the physical body of the robot - they can be external sensors that the robot has access to, like a camera on a building
- Similarly, not all of the processing has to take place on the robot - it can send data to the "cloud", process that data there, keep part of its internal representation there, and even have new code downloaded from the cloud
- All of the code/algorithms are written and tested[98] by a human, either directly (by writing the code) or indirectly, by choosing the training and evaluation data set, machine learning algorithm, and reward functions used to generate the code[99].

---

[98] Hopefully tested, and hopefully tested in as many environments/situations as possible. But it's hard - and time consuming - to conduct in-depth, extensive testing.

[99] Machine learning/AI is just CODE. The code computes very complex functions that map input sensor data to internal representation - numbers and symbols - or directly to actions, but in the end, it's just code. See https://en.wikipedia.org/wiki/Explainable_artificial_intelligence for a defnition of Explainable AI.

- Tele-operation enables an outside human operator to by-pass the robot's decision making (or change it – this is often called "human in the loop")[100]. This can be done at several levels of abstraction from manually driving the robot with a joystick to simply picking a direction for the robot to (autonomously) navigate to.

Each component of the robot - sensors, internal state and processing, actuators/robot physical shape - determine not only what tasks a robot can (and cannot) do, but also how well it can perform. To illustrate what this looks like in practice, we take a simple task that a sidewalk robot needs to perform: don't run into a human. We can break this down into sub-tasks that each component needs to perform. The seemingly simple task quickly becomes complex, and as is always the case in design, there are trade-offs between time/cost/energy and functionality/safety/reliability. We'll start by considering how a robot can avoid running into people, and show just how complicated this process is. Then we try to make explicit as many of the trade-offs as possible that exist between regulation and design, just for this one basic task. You can then amplify the scope of our observations when you consider the considerably broader range of tasks that a sidewalk robot needs to successfully perform in the process of moving from point A to point B through spaces that are meant for, and occupied by, a range of people and things. This is the sort of analysis we suggest needs to be more explicit in law-making, as we explore in more detail later when we discuss the implications of some of the legislative decisions that states have made so far in regulating sidewalk robots.

Take a moment and think through how you, as a human, avoid running into people while walking down the sidewalk. Do you actively record every person and make a mental list for each person of who they are, what they're wearing (so you can recognize them again in the next second), where they are relative to you, how fast they're moving? Do you make eye contact or exchange non-verbal information (such as slightly shifting your direction of travel or looking to the side) with the people? Do you do this for everyone, or if not, how do you pick the people you track and/or make eye contact with? How do you know if a person is standing still or moving? How do you know when a person is going to start (or stop) moving or change direction or speed? How do you know how close a person is to you and how fast they're moving and if their path will intersect yours? Do you notice (and if so, how?) if a person is on a scooter, pushing a bicycle or stroller, roller skating, using a cane? How close are you willing to let the person get to you, and does this change if it's a crowded

---

[100] There is a growing body of literature [https://hai.stanford.edu/news/humans-loop-design-interactive-ai-systems] that seeks to understand when and where humans can successfully step in and deal with cases/problems the robot's autonomous systems can't handle – so-called "human in the loop".

environment? Can you tell the difference between a poster with a human on it and a real human?

In broad strokes you, as a human, take in sensory input (vision - identifying people and where they are with respect to you, sound - footfalls, bicycle wheels on pavement) and map this into an internal representation (an estimate of where people are and what direction they're moving, probably filtered by if they are close enough to you to matter) and use that to continuously adjust your walking speed and direction as well as your eye gaze, head orientation, and facial expression. There are very complex operations going on in your head, learned through years of walking around in the world: You can track/follow a person by glancing at them every few seconds, recognizing that it's the same person even though they've changed pose, the lighting has changed, and their clothing is a different shape. You can reliably estimate - again from a half second glance - where that person is going to be in a few seconds, regardless of if they're on a skateboard, pushing a stroller, or their walking speed. You can recognize if a person has seen you (and knows you saw them) and has adjusted their path to avoid a collision.

Hopefully, at this point we've convinced you that the task of moving down the sidewalk without running into a person is really, really hard and that we, as humans, make a lot of assumptions about how other people move and that we rely on other people to contribute to accomplishing that task. We also have very little introspection about *how* we actually do this reasoning. So what does this look like on a robot?

Similar to a human, a robot has a set of sensors that measure the physical world in some way. This sensor data is then processed into a form that matches the robot's internal state. This internal state, plus a representation of the desired goal, are used to determine how to activate/move the robot. At a superficial level the human and robot seem very similar, but the sensors and internal state are so different that it's dangerous to apply human terms to what the robot is doing.[101] The following are three sensor/state/actuator combinations and how they model this task that illustrate just how different the solutions can be. We sought to identify the trade-offs for regulation and the practical operation of a sidewalk device for each of these combinations. We hope these give some insight into the range of possible requirements for sidewalk robots, and what these requirements would mean for the people who have to encounter them in shared public spaces, like sidewalks, crosswalks, and curb cuts.

---

[101] https://www.brookings.edu/techstream/the-danger-of-anthropomorphic-language-in-robotic-ai-systems/

Scenario one: Minimal sensing and internal representation

(Trade-off: easier to produce, less expensive, but less sophisticated at the task of not running into people.)

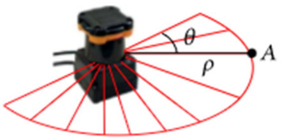

In this scenario, the robot has two object-detection **sensors** - a laser scanner that scans at knee height, four bump sensors located at the front, back, and sides, and two "cliff" sensors, one at the front and one at the back (for detecting the edges of the sidewalks). The laser sensor returns 180 values every 1/30th of a second, representing the distance to the closest object in an arc in front of the robot [need pic]. The laser only works out to 2 meters, so past that it can't see anything. The bump sensor is triggered when the robot actually bumps into something. The "cliff" sensor is a simple time of flight sensor that detects the distance to the ground in front of the wheels.

The robot also has three localization **sensors** - a GPS, an Inertial Measurement Unit (IMU, that returns the robot's linear and angular accelerations, which can be used to calculate velocities and positions), and odometry from the wheels (reporting how much, and in which direction, they have turned). Similar to how your phone works, these have enough information that the robot knows where it is within 1-2 feet, and what direction it is pointed in within 5-10 degrees. The robot also has an internal map of the sidewalks in the area.

The robot has **four wheels and drives like a car**; it can go faster and slower, and turn, but it can't move sideways without doing a parallel parking maneuver.

Internal state: Where am I? The robot has a map of the area it drives around in, and can estimate where it is in the map. This is part of its internal state. As part of the training process, the robot was driven around the space without any people in it; this was used to build up a very accurate map of what obstacles the laser bounces off of around the sidewalks (eg, bushes, trees, park benches - things that are at knee height). Using this, the robot can estimate where it is (provided no-one moves anything and there aren't a lot of people in the space so it can still see parts of those objects). Note that this map is entirely at knee height - the robot cannot "see" anything above or below that height.

Internal state: What's in front of me? Every 1/30th of a second the robot takes a laser scan. This laser scan is compared to its internal state to determine what the distance values **should** be. If the distance value is shorter than expected, it tentatively places an object (cylinder) at that location. If it gets enough consistent readings, then the probability of that object actually existing (and not just a sensor glitch) goes up. Through a bit of math, a few more sensor readings, and knowing how far it has traveled, it can determine that there is an object at that location and what speed and direction it is heading (making some assumptions about how fast objects move). In summary, the internal state is that there are currently N objects, each with a position, heading, and likelihood of actually being there (since sensors are inherently inaccurate).

Decision making: Which way should I go? The robot has a target destination on the map (given to it by an external source, such as a delivery app). The robot plans a path from its current location to the target using the map (think Google maps or the equivalent).

Decision making: How fast should I move and in what direction? The robot has two potential decision-making components. The first uses the internal state (a set of locations of objects and velocity vectors) to calculate potential collisions. It can then calculate changes in heading and/or speed that (potentially) reduce the risk of collision. This algorithm has a "knob" that can be turned to trade-off speed of travel for safety.

The second decision-making component has no real "intelligence" - it just stops the robot if one of the bump sensors is triggered or if the cliff sensor detects that the distance to the ground changes (i.e., the edge of a sidewalk or a small object in front of the wheels). This is an example of a redundant back-up system.

Human-level tasks the robot can (and cannot) do:
Because the robot only has a distance sensor, it can't distinguish one object from another - or tell the difference between something that is stationary because it can't move (like a park bench) or an object that just isn't moving (like a person standing and talking to someone). So it cannot do any task (such as moving more slowly when children are present, or going closer to a backpack placed on the ground in order to avoid getting too close to a person) that requires knowing what is around it.

The robot laser sensor only covers the area in front of the robot and the robot cannot turn in place[102]. Therefore, it can't back-up safely. So if something is placed directly in front of it the only option it has is to come to a stop and wait. Or, back-up very, very slowly and rely

---

[102] You might be thinking "that's a terrible design". Well, one of the main research robotics platforms - the PR2 from Willow Garage [in 2010 https://robots.ieee.org/robots/pr2/] - weighed 300 pounds, only had forward facing cameras, and NO bump sensors.

on the bump sensor to detect if it runs into anything[103]. In summary, the robot cannot (safely) perform more advanced avoidance maneuvers such as “back up and go around”.

If the robot stores not only the map, but also past path travel times, it can alter its path planning to account for previous times when it got “stuck” (similar to avoiding a particular intersection at 8am because of commute traffic).

Privacy and security: The robot can collect usage statistics - how “occupied” a given sidewalk is at a particular time, but it cannot collect identifiable information (who, or what, was in the space). Although the robot is reliant on an externally-provided map and GPS data, failure or “hacking” of either of these will not cause the robot to actively run into anything (or off of the sidewalk[104]) because the internal representation of objects depends primarily on the laser.

On the robot versus “in the cloud” data storage and computation: The majority of the decision making needs to happen on the robot in order for it to be fast enough. The updating and storage of maps would primarily happen in the cloud, especially if multiple robots were navigating in the same space - the map could be updated with each robot’s data, and the new map downloaded to the robot as needed. However, the map updating could be accomplished on the robot without need of cloud computing resources.

Scenario two: Adding a camera and mobility

In this scenario we take the robot from the previous scenario and add a forward-facing camera and the ability to spin in place by enabling the wheels to actuate independently.

Trade-offs: safer functionality, but more data collection & privacy concerns, cost

Sensor processing/fusion: An additional processing unit is added that “fuses” the camera data with the laser scan data, enabling object detection and labeling. Essentially, the camera image data is processed to create bounding boxes with possible labels (todo Figure) and the laser data adds distance information. This processing happens on the robot and is trained before deployment with a standard set of objects (cars, people, bicycles, etc) found in a sidewalk environment. Similar to the map creating above, the robot is driven around the

---

[103] Which leads to the amusing situation in a crowded environment (or a bump sensor failure) of the robot just… sitting in the middle of the sidewalk because it can’t safely move forward or backward without running into something. And that something may be… another robot.

[104] It could, potentially, drive down a ramp - but the fact that it is on the ramp would be detected by the internal IMU sensor.

space for a few days and the labeling data manually hand-tuned for the environment, so that the robot has a reliable map not only of expected distance to objects, but example image representations of those objects.

Internal state: Where am I? This is very similar to the previous scenario, except the robot also has additional information from the camera to help it localize even in the presence of obstacles. This is because it can match image data (the tops of trees, the sidewalk) to determine if the image is the "same" as the image data it saw before. One caveat: image data changes a lot, based on time of day, clouds, fall versus spring, etc. It is, however, more "unique" in that very few places look like other places (unlike laser data).

Internal state: What is in front of me? This is very similar to the previous scenario, except the robot also includes an image (with label) of the object, and the past location(s) of the object. Objects can be detected in both the laser scan AND the image, enabling detection of objects that the laser scan alone cannot detect (such as a person standing behind a hedge).

Decision making: Which way should I go? This is the same as the previous scenario.

Decision making: How fast should I go and in what direction? Similar to the previous scenario, the robot decides at every time step if it needs to change its speed or the direction it is heading based on minimizing potential collisions. By keeping information about where the object has been, it can better estimate where the object is going and, therefore, move faster and closer to objects.

Human-level tasks the robot can (and cannot) do:
Depending on the accuracy and reliability of the object labeling, the robot can adjust its internal representation of expected speed, motion, and potential damage during a collision on a per-object basis. This would not fundamentally change the ability of the robot to avoid running into people (or objects), but it would enable the robot to dynamically adjust its speed/risk trade-off based on what *types* of objects are around it. The programmer, in this case, would provide, eg, a preferred risk/speed trade-off for each type of object.

Tasks that involve object detection are do-able with this type of robot. These include 1) non-identifiable monitoring and detection tasks - how many people on the sidewalk or objects left on the sidewalk (like backpacks), 2) object/people identification - keeping a list of the unique faces or licenses seen over a period of time, how often a car is parked in the same location, 3) surveillance - keeping recorded video, tagged with location and time.

With the ability to spin in place, and a more sophisticated labeling of objects, the robot can plan maneuvers to, for example, back up and go around a knot of people that have clustered in front of it. The sensor range of the robot is still, however, limited to the immediate vicinity around the robot. This means it cannot plan a more sophisticated avoidance solution, such as knowing that a particular sidewalk is crowded so it should choose another route.

Privacy and security: The robot is, essentially, a mobile recording unit, so (with sufficient bandwidth or memory storage) it is capable of recording everything that happens in front of it. With access to facial recognition, it is possible to keep persistent records of who the robot encountered and when. The robot is also able to track license plates [cite police work].

Object labeling - and determining speed and heading based on object type - can result in overly risky speeds OR slowing overall delivery times with overly conservative speeds. This can happen both because the programmers make those decisions and because object detection may have biased failures. As a simple example, consider a person on a skateboard. If the object detection algorithm fails to detect the skateboard then its estimate of speed of travel is far too low.

On the robot versus "in the cloud" data storage and computation: Similar to the previous scenario, most of the computation for avoiding people can happen on the robot, once the map is downloaded. Shared map updates would happen in the cloud.

Scenario three: Adding external sensors

In this scenario, the robot itself is the same, but the environment has been modified to have cameras with depth sensors that cover the sidewalks.

Trade-offs: greater sophistication, safety, efficiency in the robots but infrastructural investments and cost (and the bureaucratic challenges that come with making them); also important considerations and trade-offs around whose interests get prioritized in infrastructural investments and modifications.

Sensor processing/fusion: There is a central system that monitors the external cameras and updates, every second, a map that has estimates of objects with labels and their locations (similar to the internal state from scenario two). The robot can access this map at any time, as well as the images from the cameras. The cameras are calibrated, so the robot knows which pixels in which cameras correspond to the locations on the maps.

Internal state: Where am I? As before, the robot keeps track of its position and orientation within the map using its internal sensors. It uses the external camera system to cross-validate its location, and continuously broadcasts its location to the external map.

Internal state: What is in front of me? As before, the robot keeps a list of nearby objects and estimates of their speed and direction of travel. It also uses the external map to keep an estimate of how densely occupied the sidewalks are for all of the paths it might take.

Decision making: Which way should I go? The addition of the external camera sensors makes it possible for the robot to plan a path that avoids any *current* congestion (similar to a map updating your route based on traffic).

Decision making: How fast should I go and in what direction? As before, local decisions about how to change position and direction are based on the local data. However, if the robot is not making progress it can now take advantage of the sidewalk density map to re-route through less densely-populated sidewalks. It can also use information about object/people density just outside of its own sensor range to pre-emptively adjust its path.

Human-level tasks the robot can (and cannot) do: The existence of a centralized map with locations of all objects (including robots) enables multi-robot coordination - essentially, account for other robot's plans (and objects in the space) to minimize congestion and avoid robots trying to pass through the same places.

Privacy and Security: This scenario extends the privacy and security issues of the previous scenario to everywhere the external cameras are pointed – similar to any camera-based surveilance system. Unlike existing systems, however, the camera on the robot is mobile and could be driven to a specific viewpoint in order to clarify data captured by an external camera

On the robot versus "in the cloud" data storage and computation: All of the external camera components are, by definition, in the cloud.

Scenario summary: What can I learn from this?

These three scenarios represent a pretty typical sequence of increasing functionality trade-off with increasing privacy and data security risks. There is also a tendancy to replace simple technology (bump sensors) with more elaborate approaches (computer vision-based) because it's "cool".

Sugar and spice: Components of a robot

There are a lot of ways to put a robot together. As a general rule, increasing the type of sensors available to the robot, increasing the complexity of the sensor processing, and increasing the complexity of the internal state representation will all increase the robot's capabilities to perform more complex decision making. This comes at a cost, of course.

- **Increasing material cost:** The components themselves - both the sensors and the computational hardware needed to process the data – increase with functionality.
- **Increasing runtime costs:** Whether or not the computation is on the robot or the cloud, this computation uses electricity. On-board computation can shorten the time the robot can deploy before needing a re-charge, while remote computation increases bandwidth needs and adds computation delays (which, in turn, can mean delays in, eg, breaking for obstacles).
- **Increasing data collection needs and risk of unexpected behavior:** More (and more complex sensors) increases the variation in the types of input the robot will see. This means more examples are needed for both for training and testing, and there is a bigger risk of missing a novel sensor combination. These novel combinations are what cause unexpected behavior.
- **Increasing privacy risks:** Leaving aside deliberate surveillance, more sophisticated decision-making means more specific and detailed data storage (eg, adding object labels, tracking when and where the sidewalks have more people).
- **Increasing security risks:** Complex sensors and sophisticated data processing capture more detailed information and require computational resources (both data and processing) that happen "in the cloud" because the robot's internal storage and processing isn't sufficient. As has been demonstrated recently[105] none of this data is truly secure.

In the following sections we'll break down the components of a robot into sensors, form factor, and internal state. We'll discuss functionality versus privacy in the context of the types of tasks a sidewalk delivery robot needs to perform. *The discussion for a different task would be very different, even with the same sensors/robot.*

---

[105] 2021 has seen a plethora of data breaches and ransomware.. https://www.crn.com/slide-shows/security/the-10-biggest-data-breaches-of-2021-so-far-

In this discussion we will largely ignore malicious attacks - hacking camera data, deliberately damaging robot-specific markers. This section is meant to introduce some nuance to the consideration of potential and unintended consequences that can come with regulations aimed at, for example, increasing the "safety" of sidewalk robots where safety is defined as not running into people.

Sensors and sensor processing

Distance sensors

Examples include LiDAR, RADAR, ultrasonic, infrared, depth (RGBD) and stereo cameras.

Functionality

All of these sensors push energy into the world and measure that energy as it comes back. This provides reliable measures of how far away objects and the environment are in the direction the sensor is pointed. Their limitation is that they can only measure distance in the direction(s) they are pointed - and the more directions they sample in the longer the sampling time.

All depth sensors have objects they cannot detect - lasers pass through windows, miss skinny objects, and can be confused by highly reflective surfaces. RGBD and stereo sensors can miss objects that are too close, too far, or have a lot of surface color. However, despite these problems, depth sensors are indispensable for providing the information robots need to safely move around the world without running into things.

RGBD and Stereo cameras trade-off sampling density for accuracy - they sample in the entire field of view at once, but they tend to have lower accuracy and more failures (where they just don't detect anything or their depth estimates are way off).

Privacy-security-functionality trade-offs

Distance sensors are more effective for *class level* descriptions over *instance* ones[106] - i.e., they might be able to determine that there is a car in that direction, and possibly the make and model - but not specifics that *identify* the specific car, such as bumper stickers, color, if there is a car seat in the car, etc. The exception to this is RGBD and Stereo cameras, which match the distance measures to visually identifiable information.

---

[106] Class or object level descriptions means you know the type or label of the object (cup, car, person) but not if it's a specific car or book.

Distance measures can be used as-is for either localization (does this set of distance measures look like ones I saw before?), object avoidance (stop if something is in front of the robot, stay centered in the sidewalk), and map building. Additional processing can be used to identify specific types of objects, or separate moving objects from static environmental features like walls and sign posts. Using temporal data - and information about how the robot is moving - individual sensor readings can be linked together to track objects as they move around and/or track how the robot moves. Object tracking and detection - while increasingly more reliable - is still prone to missing or mis-identifying objects and losing tracking.

Identifying static environmental features (like corners and doors of building) can improve a robot's localization, which in turn can allow a robot to move faster because it knows where environmental hazards (such as edges of sidewalks, walls, upcoming corners) are. Improved localization can also improve identification of (and, potentially, avoidance of) non-stationary objects. In general, these improved environmental maps and class-level object detection algorithms have minimal privacy concerns *in public spaces*[107]. However, as we'll see in the Internal State discussion, adding in temporal, long-term tracking of objects does have privacy implications.

With the addition of camera data (which can be spatially "overlayed" on top of the distance data) objects/people/places can be uniquely and individually identified. This individual object identification (knowing who a person is or what type of car) adds minimal benefit for robot localization, navigation, and object avoidance. Obviously, the increased risk to privacy drastically increases if the robot is identifying specific cars and people.

Time-expense-functionality trade-offs

A hidden cost of increasing sensor complexity (and functionality through additional processing) is speed of response. A time-of-flight sensor placed in front of the robot that receives data 120 times a second is going to detect a child jumping in front of the robot in a fraction of a second, whereas a camera with a depth sensor taking pictures 15 times a second, with additional processing time of a second, may not detect that child for a full second. A path planning algorithm can afford this delay, whereas a safety algorithm cannot.

---

[107] Private spaces are another matter - although having a Roomba [https://homesupport.irobot.com/s/article/29814] build a complete map of your house may help with its cleaning strategy, the company now has a complete floor plan of your house, including the amount of furniture (and potentially, the type) in your house.

Camera/image sensors

Cameras can either be placed on the robot, or in the environment the robot operates in.

Functionality

Robotics often uses omni-directional, fish-eye, or wide field of view cameras to increase the amount of the world the robot can see. Cameras can be fixed in position, or mounted on a gimbal, allowing them to pan across the scene. Cameras also have an image rate - how fast they can take pictures and image resolution, which determines the size of recognizable features[108]. For an object to be recognized in the image it must have a sufficient number of pixels and be in the camera's field of view.

Object detection: Image processing algorithms based on machine learning essentially draw boxes around parts of the image and label them with object labels. Functionally, they work by comparing the pixels in the image to a huge number of images it has been trained with (which are labeled with boxes) and determining which ones are similar. They might further "cut-out" the pixels belonging to the object (not just put a box around it) but usually the algorithms just return boxes.

It does not make sense to talk about robots detecting objects in images without also talking about the data set used to *train* the algorithm. The closer the set of training images is to the types of images seen in practice - and with a similar distribution of objects, views, and varied lighting conditions - the more likely the object is to be recognized correctly. Because these algorithms are, essentially, comparing numbers to numbers, their failure modes are hard to predict, particularly for images that are not like any image in the training set.

Object localization: By comparing images between cameras, or the same camera at two times (stereo vision), or combining cameras with a depth sensor, it is also possible to place the identified object in space. This can be applied to both the environment (walls, sign posts) or objects moving in the environment.

Cameras make it very easy to do instance-based object detection (a specific person or car) but they can also operate in class-based object detection or just keep detected depth, similar to a depth sensor.

---

[108] The TV show CSI aside, you need a sufficient number of pixels to make out an object/details. This is a function of the field of view and the image resolution. The wider the field of view, the more pixels are needed to take a picture of the same object at the same resolution.

Privacy-security-functionality trade-offs

As in the previous section, improved localization and class-level object detection can improve the robot's navigation and object-avoidance capabilities. However, the trade-off here is not as clear cut, both because of the processing speed/accuracy of the data and the nature of the data itself (images). It is very difficult to scrub identifying data out of images[109] and difficult to do object-tracking or localization without recording (at least temporarily) that identifying data in order to find it again in the next image(s).

One more subtle point is that the distance data is *derived* from the image data - it is computed by, essentially, comparing pixels from one image to another and doing a bit of math. True distance sensors actually measure the empty space. This makes depth data from images a lot less reliable and prone to error (in addition to being slower to acquire).

Cameras placed in the environment (instead of on the robot) present an interesting trade-off both in terms of privacy and functionality. Environment-based cameras can provide information about places the robot cannot directly detect with its own sensors, which is advantageous for path planning. Unlike robots, they also record information all of the time. However, the field of view of these cameras can also be made public and be clearly delineated or set to only record a specific portion of the environment. Robot-based cameras, on the other hand, can be driven to places to explicitly record images from private places.

Time-expense-functionality trade-off

Environmental cameras are an amortized cost, both in terms of material costs and runtime costs. In an environment that has multiple robots this shared computation can help reduce runtime costs for the individual robots (not to mention, they can be plugged in and don't require a battery).

Identifiers

Although not strictly a sensor, there are a number of ways for objects/people/places in the environment to carry unique ids. These include: Bluetooth signals, RFID tags (limited distance range, but don't require line of sign), QR codes or labels (license plates, signs, door plaques), phones (GPS, Bluetooth, cell-phone tower-enabled tracking), and clothing such as

---

[109] Google street view started with capturing unintended funny views, but turned into an industry all its own: https://www.cnet.com/pictures/crazy-images-caught-on-google-street-view/3/

reflective safety jackets. GPS[110], itself, can be thought of as a unique identifier (the location of the robot with respect to the satellite).

Functionality

These identifiers can be used either to identify which objects are in the environment and/or, to a limited extent, *where* that identifier is. If the identifier is tied to a specific location - like a street corner or a door - it can be used to localize the robot

Privacy-security-functionality trade-off

Markers in the environment can greatly improve a robot's localization capabilities[111] with very little impact on privacy (at least in public spaces). De-identified markers - recording that there is a person with a phone, but not which person or phone - can improve a robot's path planning.

Object-type tags - particularly ones that can broadcast their position - could improve safety for objects the robot should absolutely avoid hitting, or help the robot to avoid a situation where the sidewalk is not wide enough for both the robot and the object. Although using object type (stroller, scooter) - can reduce privacy risks, it does not eliminate them if information can be linked to other sources.[112]

Time-expense-functionality trade-off

Using tags that are easily identifiable - even if they are image-based - greatly reduces the amount of computation needed to find and recognize them. They also tend to be more reliable, since they are designed to have unique visual features and be easy to find.

Form factor

The physical appearance of the robot - and physical indicators like turn signals, visible cameras and *how* the robot moves - are how the robot communicates with the people in the

---

[110] It's tempting to think GPS "solves" the localization problem, but in practice, there are several things that can go wrong. First, the built environment changes signal strength and accuracy, even within the same block [https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0219890]. Second, GPS is deliberately set by the US government to have approximately 16 ft accuracy. In practice, most mapping devices use a combination of the street map plus IMU data to produce better localization. You may not veer 6 feet to the left into on-coming traffic when the GPS lock on your phone fails, but a robot might.

[111] Amazon's warehouse Kiva robots

[112] Credit card receipts plus cell-phone data, for example.

environment. While regulation has focused on weight, size and speed maximums,[113] these other elements of the robot's physical design are just as important. Anecdotally, one of the biggest problems with Starship robots is that the people around them have no idea what they're doing, especially when they're stopped at an intersection. Are they waiting for traffic to clear? For you to cross? We'll look at each of the form factor elements in turn.

Size, overall shape, and visibility

Do you make your robot short, and wide? Or tall and narrow? Should it look sleek and light, like a racing bicycle? Or stocky and slow moving? Should it look like hard metal or like a soft beach ball? Is it visible or eye-catching, particularly at night? Humans are actually very good at estimating overall weight and material properties from visuals alone; hence the large amount of money spent on designing car bodies. There are three primary factors that go into the robot's overall size and shape:

- The environment the robot is in may dictate limits on overall width or height, such as keeping the robot less than half the width of the sidewalk or less than a certain height so it will fit in a tunnel.
- Carrying capacity and the need to carry electronics (all those sensors and processors take up space, and battery life is limited by the size of the battery). This will tend to push overall size as big as possible for most applications. The counter to this is that tall robots are more likely to tip over, and bigger robots cost more in materials.
- Form matching function: A hot-food delivery robot (sleek, fast, not cuddly) might look different than a child-minding robot (short, slow, and cuddly) or a sidewalk cleaning robot (slow and sturdy). Similarly, graphic design and the robot's shape can be used to cue people into what the robot is *for*, and so how they should move around it (or get out of its way).

Although there are practical considerations constraining the overall size and shape of the robot, the external form factor will have the biggest impact on safety and functionality in the context of sidewalk robots. Simply put, sidewalk robots will need to be around the size and shape of a rolling suitcase or stroller (or the size of a bicycle with a trailer if operating in the bike lane). Past that, the robots need to be visible - but not distracting - and clearly communicate their purpose so people can move around them comfortably.

Static form factor elements, like lights, flags, obvious cameras, and bumpers, are also key methods for communicating (flag - makes the robot more visible, camera/wheel orientation - communicate that it has a camera, which direction the robot will move in, bumpers - it may

---

[113] See first section.

run into objects). One issue that already has arisen is contact information - is an email or website sufficient? A logo? Should robots be required to have a phone number[114]? Another point of concern is visibility at night – are reflective stripes enough? What about head and break lights?

These are all factors that might be of interest to law-makers who are concerned with integrating sidewalk robots into public spaces in ways that are safe and equitable. But at the same time, these are all factors that would be very difficult to effectively regulate with any specificity.

Weight, speed, maneuverability, and agility

Existing legislation has tended to focus on weight separate from speed, with, presumably, the intention of keeping them "safe".[115] We argue that these elements - plus maneuverability and agility - should be treated together (along with requiring basic collision sensors). A lightweight robot traveling faster than a slower, heavier one can impact with the same force. If a robot can't maneuver around obstacles or safely roll off the curb then it risks becoming a hazard. A smaller, less maneuverable robot will result in fewer safety concerns than a bigger, more maneuverable one, especially in a crowded environment, simply because it won't hurt anything if it runs into it.

Maneuverability: We have deliberately side-stepped whether the robot has wheels or legs, how many wheels, how they turn, etc. This is because maneuverability is a function of the physical design (number of wheels, legs, etc), how they are actuated (speed of the motors, how far the wheels can turn, turning radius, etc), and the sensor package (how quickly, and at what distance, the robot can detect an obstacle and how accurately it can detect that obstacle).

Agility: We use this to refer to the robot's ability to handle curbs, rolling off the sidewalk into the grass, sidewalks rucked up by tree roots and pot holes, etc. Agility is mostly a matter of kinematic design for wheeled robots; for walking robots, sensors and control strategy play a much bigger role.

---

[114] https://www.bgfalconmedia.com/news/starship-vs-whips-robots-pose-potential-risk-on-streets/article_f31331b2-91af-11eb-af2b-d7feb69985a3.html

[115] In reality, some of the legislation may be designed to make their competitor's robots too heavy or fast.

As in the previous section, there are multiple factors at work; in general, increasing functionality (through more sensors, processing, or motion capabilities) tends to increase weight. Heavier robots, moving fast, are less maneuverable (physics). Unlike previous sections, however, increasing maneuverability and agility might increase material costs, but doesn't tend to increase run-time costs (and, in fact, may reduce run-time costs through passive physics).

Factors driving design:

- Operational environment - how smooth and wide are the sidewalks, is it hilly, how crowded are the sidewalks, what is the terrain on the side of the sidewalks, are there curb cuts, are there curbs, are there fences/buildings/hedges etc. What are the implications if the robot gets stuck/breaks down? The more challenging the environment, the more maneuverability and agility is needed.
- Robot design and maintenance constraints - What is the robot carrying/delivering? Is it a general-purpose robot that is expected to be deployed in a wide set of environments, or is it a specialist? What is the desired maintenance/replacement/break-down trade-off (cheap robot that is replaced frequently[116] or expensive robot that is expected to run for a long time between repairs)?

Here we give a potential list of descriptive functionality measurements that should be used instead of prescriptive descriptions like weight.

- Stopping distance at different speeds: This is a measure of how quickly (and reliably) the robot can detect an obstacle and then stop.
- Maximum force in the event of a collision: This is a function of the weight and speed and what the robot is built out of (think car bumpers, which are designed to absorb force).
- Movement ranges/turning radius/passage width [need a better term, maybe split into two]. Basically, what directions can the robot move in, given a specific environment? This is a function of the mechanism - i.e., a car can move forward and turn, but can't move sideways, sensor accuracy - a robot may be able to physically make a turn/fit through a space, but because the sensors are inaccurate (it doesn't quite know where the walls are) it can't fit through the space without bumping into the walls, and control (particularly for walking robots that have to consider where the feet go while moving).
- Robustness to terrain: What obstacles (like potholes) can the robot safely roll over?

---

[116] Because robots can be largely interchangeable, it may make more sense to just have a few extra around and swap out when needed.

Although descriptive capabilities are more time-consuming to measure (and require a suite of tests, rather than a simple weight measurement), they more accurately measure the robot's behavior in practice. Law makers need to be attuned to what measures will actually achieve their stated goal, and whether regulating those measures is actually practical. Another trade-off that arises here is enforcement – both its practicality and its cost. Laws that do a better job of regulating the robot's behaviour in practice also become more difficult to enforce (e.g. weighing a robot to ensure it meets a simple weight requirement vs running a series of tests to certify that a system operates the way it's required to).

State indicators - external communication

In addition to static form factor elements, a robot can have *active* state indicators. A very simple indicator is wheels - which way the wheels are pointed communicate which direction the robot will move in, and turning wheels indicate a change of direction. A camera mounted on a movable gimbal can communicate what the robot is looking at. Brake lights and turn signals can indicate the robot is stopping or turning. Audio is another state indicator, ranging from a simple beep when backing up to voice.

Increasing the complexity of the state indicators can result in more complex communication - but also more confusion. Does a blinking green light mean the robot is thinking, signalling for you to go, or that it is turned on and waiting for commands? There is a reason that car state indicators are still limited to just turn signals and brake lights[117]. In environments where people are exposed to the robots regularly, and the robots perform the same actions in the same situations, people will learn those signals fairly quickly - *if they can be interpreted largely unconsciously*. This is the study of *affordances*, and, in summary, it should not be assumed that just putting some state indicator on the robot will actually communicate the desired intent.

An indirect state indicator is the actual movement of the robot - does it move slowly, smoothly and in a straight line? Does it stop and start? Does it change direction quickly or slowly? All of these contribute to whether or not people in the environment can (largely unconsciously) understand where the robot is going and what it is doing, which in turn helps with negotiating passing each other on the sidewalk, moving in a crowd, and yielding right of way when necessary.

---

[117] Person to person communication such as eye gaze - and subtle things like wheels turning and changing speed - play a huge role in our ability to interpret what another car is going to do.

Similar to maneuverability above, it is possible to define descriptive measures of the effectiveness of state indicators. Unlike maneuverability measures, however, the appropriate use of these measures is not *safety* but *resource utilization*. A robot that clearly and smoothly moves through space may enable more people plus robots to use that space - but it won't, in and of itself, prevent serious accidents. It may, however, reduce the overall number of times that minor accidents and incidences happen.

Internal state

What goes on "under the hood" of a robot is largely proprietary - but there are some broad characterizations that can be made. Specifically:

- Internal versus external (or cloud-based): Internal storage is far faster to access, but limited in space. In general, data that is not time sensitive and more global (like maps and congestion) tends to be stored externally. Internal data that captures the local state (where objects are, local features of the environment) are kept on the robot. There is both a privacy and a safety trade-off, with local data storage both reducing privacy risks and increasing safety. External, shared storage, however, can increase overall efficiency by reducing congestion and enabling more informed path planning.
- Long term and persistent state versus just-in-time state: A robot can be entirely reactive - not storing any state over time, store just a limited temporal window (tracking objects and people in the immediate environment), or keep long-term statistics on objects and people it encounters[118]. Longer-term state requires more computational power to process, compute and reason about, but can bring path efficiency benefits. It can also raise concerns around privacy and data collection.
- De-identified or identifiable data: There are some simple techniques - blurring, reduced dimensionality, aggregated data[119] - that can be used to reduce the privacy risks associated with data storage.
- Shared versus private data: Any part of the robot's internal state - from the robot's position to map information to locations of objects in the environment - can be shared.

It is very tempting to think about robot internal state as if it were the robot's "memory", similar to what a human might remember from a walk. It isn't. A human might remember a few key events or objects that stood out - a friend, a garbage can in the sidewalk, a blooming flower, a more-crowded than usual sidewalk - but not every object and person encountered.

---

[118] This is very similar to how map applications work now - they can provide no traffic information, current traffic information, or a guess, based on historical data, of what the traffic will be like at a specific time.

[119] Averting Robot Eyes, Maryland Law Review Vol 76, pg 983, 2017, Margot Kaminski et. Al.

Robots, depending on how they're programmed, can range from saving no data at all to recording when and where every unique face and license plate was encountered. Unlike humans, this data can be near-instantly shared to any interested party.

Human in the loop

The discussion has, so far, assumed that the robot is autonomous - operating without a human in the loop. There are many detailed characterizations of how humans can be "in-the-loop";[120] for sidewalk-roaming robots the most likely role of the human is supervisory (intervening when the robot gets "stuck"). It simply doesn't make financial sense for each robot to be individually controlled (tele-operated) by a dedicated human. Similarly, the potential cost of having a robot stall-out in the middle of an intersection and cause an accident means most robots will be supervised by a human - the autonomous technology just isn't reliable enough, yet, to leave them unsupervised.

The first myth to dispel is that a human can step into an on-going, time-critical situation and fix it. Fortunately, most situations a sidewalk robot might encounter are not time-critical at the split-second level. There are many situations that might require a human to step in and either directly tele-operate the robot or provide some form of decision making.

- Situational awareness prior to potentially dangerous actions such as crossing a street.
- Strategic re-planning or intervention to avoid known problem areas.
- Monitoring of robot state to detect break downs, the robot getting stuck, or deliberately obstructed.
- Communication with humans interacting with the robot.
- Direct tele-operation of the robot to navigate obstacles (eg, if the sensors are not accurate enough to safely move the robot through a narrow opening).

Although humans can bring more insight into a situation they are not a panacea for all problems a robot might encounter. In particular, humans are limited to observing the robot's world through the same sensors that the robot has.

The environment (sidewalks, external sensors)

There are a variety of environmental elements that can play a role in how a robot navigates the environment. We break these into sensors, actuators, and the physical design of public spaces.

---

[120] [https://hai.stanford.edu/news/humans-loop-design-interactive-ai-systems]

Because these are city-wide, public-space infrastructural changes, who decides which are worth doing - and where - and who pays for them are not easily resolved questions. Another challenging question is who owns the data, and who has access to it. Urban 'smart city' projects like Sidewalk Toronto signal only some of the many challenges that can come with environment regulation and change.

Sensors

External cameras or distance sensors, GPS-like devices for localization, markers embedded in the side-walk, and maps are all sources of data robots can use. In the case of cameras or distance sensors the robot could query the raw data or an external system could process the data and provide data in a variety of forms. Markers - ranging from images embedded in the sidewalk to RFID tags placed at corners - provide passive mechanisms that the robot can use to localize itself.

The environment can also be modified to directly broadcast data, such as the state of stop lights, if doors are open or closed, or if a train is on the tracks. Pressure sensors in roads, in addition to triggering stop lights, could be used to broadcast the presence of a car or bicycle.

Actuators

Stop lights - or crossings - could similarly be automated to enable a robot to trigger them.

Physical environment

There are many modifications that can be made to the physical environment to increase safety and reduce congestion. Just as there are dedicated bicycle lanes, there could be dedicated sidewalk robot lanes. Just as there are pedestrian over or underpasses, tunnels could be added under or over busy intersections or under existing bicycle lanes. Curb cuts and robot crossing markings both make it easier for a robot to navigate and make the robot more visible.

Bringing It All Together: Potential policy questions

The primary premise of this paper is that legislation in and around sidewalk robots needs to be informed by technology (what is technically possible, what can actually be accomplished with certain design regulations, and what are the trade-offs for different technical requirements in design and/or functionality). We also want to emphasize that some of that "technology" is not contained in the robot itself, but arises in how shared public spaces are built and maintained. There are fundamental questions about how, and what kind of, data can or must be collected and who has access to it. Similarly, robots occupy physical infrastructure and can cause both safety issues and congestion - how should this shared physical space be shared with robots (and even more fundamentally, should this space even be shared with robots?)? How do we ensure as much safety as possible? How do we address safety concerns that go beyond physical injury or property damage, into the realm of the safe lived experience of public space? Who pays for necessary investments, and how? Whose interests get prioritized in all of these decisions?

Another key question in all of this is how best to approach regulating: should law-makers write pro-active laws/policies, or wait and use the courts and existing laws to react when something bad happens? Proactive laws sound intuitively better until we consider the enormous range of possible actions and encounters that can be implicated by a simple instruction like, don't run into people. It would simply not be possible or effective for law-makers to preemptively regulate every possible dynamic that can arise from introducing sidewalk robots into a public space. But how should law-makers identify which proactive laws might be helpful? We don't have the answers to all of these questions, but in the sub-sections below highlight some of the nuance that needs to be considered when regulating robots to enter shared human spaces.

Data collection and privacy in public spaces

Robots are, essentially, mobile data collection devices. Add in external cameras, tracking of cell-phones, and face-tracking and it becomes easy to track individual movement. We make the argument that de-identified, occupancy-based data is a feasible compromise between increased robot capabilities and privacy risks. We also argue that legislation is needed to prevent and prohibit deliberate tracking in public spaces. Finally, a broader discussion around what type of data sidewalk robot companies might be required to share and with whom is needed.

We argue that data requirements should be put in place up-front and not left to post-facto law suits. It is very difficult to establish harm if you don't even know what data a company has recorded about you. It is also challenging to put data "back in the box" once collected. Leaving what is acceptable - and not - undefined also risks turning public places into de facto corporate profit-generating spaces. And places the burden of addressing privacy violations on the shoulders of those who have been harmed, instead of on those who stand to profit from information collection. There are several considerations that should be in the minds of legislators as they attempt to regulate sidewalk robots:

Data tracking that improves safety and reduces congestion

Local object tracking is a tool that can improve the immediate safety/safe operation of robots moving in and around complex environments. Similarly, congestion maps - either collected and shared between individual robots or collected via shared sensing in the environment – can reduce the risk and likelihood of collisions.[121] Both of these - implemented properly – can have minimal privacy risks. Implementation matters, however. The same functionality (avoiding obstacles and crowded sidewalks) can be achieved with privacy-preserving methods or ones that track individual data. The following are steps that can be taken to minimize unintentional privacy violations:

- Minimize identifiable video data, particularly data that are stored long-term in the cloud and data made available to other entities. There are two mechanisms to achieve this. The first is filtering/de-identifying[122] the raw data. This can range from a simple blurring filter to reducing the data to bounding boxes with object identifiers. The second is to process video data into these abstract internal representations - person moving from here to here - on the robot, and only upload abstract representations.
- Publish only aggregate data, and only when there is sufficient data to be aggregated. This aggregation should be both spatial and temporal. For instance, instead of recording the track of each individual person or object encountered, store an average occupancy and direction of movement for blocks of time in sections of the map. Instead of recording a single picture of a building for localization (which may have identifiable images in it[123]), record multiple pictures over multiple times of the day and remove any non-stationary objects.[124]

One other aspect of data tracking that warrants discussion is training bias and specificity. Most machine-learning algorithms require large amounts of labeled data to be able to identify objects and learn to maneuver around them. An algorithm trained on downtown New York images will probably not work very well in a leafy Chicago suburb. An algorithm trained on images of people in a business district may have trouble with kids on skateboards. From location to location, we might find differently-shaped or painted fire hydrants, fences, curb cuts, stroller types, mobility assistance devices - all of these can confuse an algorithm if

---

[121] The difference between these two is very similar to the difference between each cell phone/GPS unit in a car individually uploading its current speed and location and a traffic camera. In the former, aggregate speed data is collected in real time from multiple sources, enabling a reasonable estimate of current travel times everywhere these devices are enabled and sharing data. Traffic cameras monitor a fixed location, and some effort is needed to (indirectly) measure the speed of traffic in the observed field of view.

[122] Averting Robot Eyes, Margot Kaminski et al, https://scholar.law.colorado.edu/articles/728/

[123] In the early days of Google's street maps they unintentionally recorded individual people in many of their vies.

[124] This can actually improve most localization algorithms, which do better on average, well-distributed data.

that image data was not in the training set OR because the *distribution* of data does not match the environment[125].

While it is fairly easy to prescribe privacy measures, these more subtle algorithm issues are perhaps better addressed with systemic measures (see below).

Some potential commercial applications for data collection, beyond efficient, safe, delivery include:

- Selling advertising space on the side of the robot based on estimated crowd density.
- Selling targeted advertising based on identification, ranging from specific (individual) to demographics-based (age, race, gender, etc)
- Selling occupancy/traffic data to vendors in street fairs/open air markets

Legislators should be mindful of the ways in which data can be used for, and beyond, the purposes of improving safety and effective movement of robots through public space.

Preventing (or reducing) deliberate tracking

This is a broader policy/privacy debate, which overlaps with other discussions around images taken in public places, including CCTV, cameras on police cars that record license plates, and, of course, drones. Although a full discussion of the issues is beyond the scope of this paper, we note

- Are you licensing/legislating a sidewalk-anything robot, or a sidewalk delivery robot? Should a robot have to identify/make clear its function/task? Should monitoring robots be legislated differently than delivery ones? We argue that they should, both because their data collection modes are different and because they move through public spaces in different ways.
- Because sidewalk robots are mobile, and thus are capable of privacy-invasive behavior than static cameras. For instance, they could take pictures of the inside of your living room when you open the door to collect your package, or note the make and model of the car in front of your house. They can drive by bus stops just as the bus arrives and record people entering and leaving.
- Sidewalk delivery revenue can be used to amortize surveillance/monitoring costs. While it might not make financial sense to deploy a fleet of robots just to track people/monitor spaces, robot delivery paths could be optimized to monitor a space

---

[125] E.g. if there is a larger number of something in a data set - like cheetahs - then the algorithm is more likely to say an object is a cheetah, simply because it "got that right" a large number of times during training. If all you've ever seen is cheetahs, then a house cat looks a lot like a cheetah. [This purrfect anecdote is dedicated to Sue Glueck]

with little additional cost. This has the advantage that the robots are not perceived as monitoring, but as delivery, robots

What data should a robot company disclose?

There is both aggregate and specific data that, if shared, could improve efficiency and safety, and also be used to amortize the cost of developing and maintaining shared public spaces. In increasing order of specificity:

- Aggregated usage data: This could be total mileage along specific road segments, number of times a signal was utilized, or number of robots occupying a single road segment. This type of data could be used to inform the need for upgrades/maintenance, or to directly bill companies for usage.
- Real-time congestion or location data: This could be used for dynamic pricing, identification of unusual events, or for route planning.
- Recording of collision or obstruction events.
  - Robot-specific id tags.

It should be clear through legislation to members of the public as well as sidewalk robot companies what kinds of data might be stored and disclosed, and for what specific (purportedly socially beneficial) purpose. Sidewalk Toronto and other failed experiments also highlight the importance of transparency and public input in the regulatory process.

Safety, congestion, and equitable access

Unlike self-driving cars - which largely replace existing drivers in cars - sidewalk robots are starting to occupy public spaces that were not designed for these devices. Safety is obviously a concern - already there are reports of robots hitting cars and obstructing the use of the sidewalk for pedestrians. We noted above, for instance, Emily Ackerman's first-hand account of how sidewalk robots can create unsafe sidewalk conditions for people who use wheelchairs[126]. Beyond the obvious harm of a robot running into something, there are broader congestion-related safety concerns as robots occupy these spaces in larger numbers, potentially "pushing" pedestrians off of sidewalks to go around them. Congestion itself is an equitable-access issue; as congestion increases this will impact certain user groups (people

---

[126] https://www.nbcdfw.com/news/nbc-5-responds/nbc-5-responds-what-happens-if-youre-in-a-crash-with-a-robot/2442345/; https://www.bloomberg.com/news/articles/2019-11-19/why-tech-needs-more-designers-with-disabilities

with strollers, people with impaired mobility, parents of small children, people using the sidewalk for rest/social gathering/income earning/*etc*) first. There are also broader demographic issues around which neighborhoods are impacted and where public and private money is spent to mitigate congestion - we come back to this discussion in the later section on public infrastructure.

At this point it is not even clear what legal box sidewalk robots should be put in - are they a person? A piece of company equipment? A bicycle? What rights do they have? What rights do the pedestrians and residents have? How should accidents be handled? Very few states have addressed these questions, and those that have begun to address these questions have arguably not done so thoughtfully enough. Take for example Mok's experience discussed above in which the classification of a sidewalk robot as "not a vehicle" made recourse after a collision more, not less, difficult to ascertain.

General laws around compensation and responsibility should apply in the context of a collision or injury caused by a sidewalk robot, but some specifications in legislation will be needed to ensure those laws of general application can be applied in the event of injuries or harm. When it comes to regulating safety, any legislation should focus on defining at least two items: 1) Goal-based or descriptive safety responsibilities for operators and 2) accident-handling requirements.

Safety: Robot collisions

Robots will run into cars and people, break down in the middle of the street, and cause accidents when people try to avoid them. It will take time and experience both for people to accommodate robots and for robot companies to determine best practices for software, hardware, and maintenance schedules.[127]

Too much detailed, prescriptive legislation in these early days runs two risks. First, a list of detailed specifications will tend towards companies focusing on meeting those specifications, not necessarily the *intent* of them. At this point, we don't really know what a "safe" weight or algorithm is, let alone how the combination of those (and the environment) work together. A broader "be safe" mandate gives companies more room (and more responsibility) to make their robot safer. Second, overly detailed legislation risks making litigation more difficult for injured parties. Companies that meet various weight and speed specifications, for example, can point to the requirements and say "we met those" and therefore argue they met the standard of care owed to pedestrians in operating their sidewalk robots. The more detailed

---

[127] See Applin, *supra*.

the legislative requirements (which, as we said above, might not actually make the robot's operation safer), the more persuasively a company might try to argue this. While not a fail-safe argument, it would risk confounding the legal analysis, especially when the laws that purport to create technical standards have to be interpreted by non-technical experts in a trial or settlement context.[128] We're not suggesting that legislators ignore safety, instead the opposite – that legislators should be cautious not to regulate in such fine detail that the laws inadvertently compromise safety, and recourse where injuries do happen. Consider the below hypothetical approaches to regulation:

Example - Prescriptive legislation: The robot has to be less than 100kg, less than 0.5m wide, travel at a speed less than 5km per hour, and have the company's phone number and logo on the side. This generally reflects the current approach to regulation, with our own addition of some specifications.

Example - Goal-based legislation: A robot should be able to come to a stop fast enough to avoid damage to people and stationary objects. A robot on a sidewalk should leave sufficient room around it for a person with a stroller or assistive device to pass. In the event the robot breaks down, anyone[129] should be able to push it out of the main traffic path. The robot should be visible from a car at an intersection.

While prescriptive legislation is in many ways more straightforward to implement and enforce (e.g. a company, and law enforcement, could weigh the robot to make sure it meets the weight requirement) a 99kg robot might do just as much harm as a 101kg robot. A pedestrian may be much safer if the robot is required to stop fast enough to avoid damage, particularly if enforcement of certification or testing are involved, even if the robot weighs 150kg.

Safety involves a mix of form-factor design decisions (reflective tape, flags, turn signals), movement patterns (speed, broadcasting motion changes like braking and turning) and overall sensor to actuator design (bump sensors, object detection, breaking, turning ability, agility). Some of these (reflective tape, turn signals) could be mandated now, based on experience and research in existing literature. Others (signaling intent to cross a road, whether to drive in the center of side of the sidewalk, conditions for safely crossing an intersection, form factor relative to the size and type of sidewalk) should not be

---

[128] See e.g. Ryan v Victoria (City) in Canada.

[129] This might be a difficult standard to define in a way that can be enforced.. but we caution against standards that rely on an "average" person, which have been known to promulgate ablism and inequity (like how air bags were safety tested to "average" and so ended up being dangerous for smaller women).

preemptively specified, but allowed to evolve into a set of "best practices". The same is true for developing end-to-end safety tests/benchmarks.

One area that should not be left to litigation and the court process, and instead be considered by law-makers, is equity in design/functionality. An example of this is the ability of robots to sense, avoid, and leave space for a range of humans, animals, and devices. Companies cannot be left to design robots based on "presumed" or "default" expectations of how people will appear to sensors, running the risk of a default to biased assumptions that exclude the range of human appearance and experience in public spaces. A sensor trained on walking adult data might exclude toddlers, teens on skateboards, parents with strollers, various mobility devices (walkers, scooters, wheelchairs), guide dogs[130], shopping carts, rolling luggage, people pulling wagons, *etc*.

Another consideration that should not be left to evolve through the courts is responsibility for accidents. Right now there is no consensus or clarity on how to handle accidents and general rights and responsibilities for sidewalk robots and the companies that operate them. This leaves the general public with little recourse or expectations in the case of accidents. Should they call the police to report an accident? Who to contact at the company? Where there are insurance requirements for companies, what is the process for injured people to access compensation? Standard levels of conduct are beneficial to the industry as a whole, since the bad behavior of one company can damage the public's trust in the field.

Possible requirements or policies include:

- A company contact number to report accidents or breakdowns, located on the robot itself (already included in various current laws for sidewalk robots) – contact information then needs to lead to an effective process for recourse for it to be meaningful.
- A claims department, no-fault insurance, or other dedicated mechanism for handling injury claims and broken robots.
- A centralized repository for accumulating information on how and where robots break down or cause accidents
- Robot companies should assume that some number of their robots will fail, and have plans in place to prevent those failures from becoming hazards.

---

[130] https://techcrunch.com/2020/08/11/the-robots-occupying-our-sidewalks/

Congestion and equitable access

Congestion – i.e. too many robots occupying too much of the sidewalk - can cause specific harm (forcing a person off of the sidewalk and into the street) or more general harm (denying access through blocking sidewalks, intersections, curb-cuts, or cross-walks). This isn't just an issue for pedestrians, but also for shop owners and residents, and people who make any use of these sidewalk features (impeded for example by blocked access to entry ways). Congestion of the people variety is also a concern - it is fairly easy to avoid a single robot on an empty sidewalk, but in a crowd, a robot at knee level can become a tripping hazard.

The ability for a robot company to preemptively avoid sending robots into a congested area depends on the technology. It is fairly straightforward for a robot company to track its own robots, and avoid placing too many of them in one location. Without shared data, it is not possible for a robot company to know the locations of another company's robots or know how crowded places are, particularly places where it has no robot sensors (internal or external). Even with the robot's sensors, it can be difficult to tell the difference between a very localized crowd (a group of people on the sidewalk) from a fully-congested sidewalk.

For technological reasons, overall regulation of congestion might be best done through some form of congestion pricing (see the below section). A strict rule that requires robot companies to design and deploy their robots so that people have right of way is not really feasible from a technological standpoint, though has been common within the current sidewalk robot laws. In fact, this kind of requirement could create *more* congestion where the robot ceases to move when faced with a crowd of people. However, it is reasonable to require robot companies to make a good-faith effort to yield right of way, whether it's to a single person or to avoid sending too many robots to an already congested space. Explicitly defining a requirement for robots to be able to make a good-faith effort to give way would both push robot companies to consider access in their design and provide the general public with a way to litigate egregious behavior.

Tele-operation, monitoring, and breakdowns

It is tempting to push all robot safety concerns off to tele-operation or monitoring, but in reality, this is about as effective as click-through privacy statements on web-sites and apps. Some terminology first:

- Tele-operation: The human sees the sensor data from the robot and directly controls the wheels (similar to steering a car)

- Shared-autonomy: The human sees the sensor data and provides indirect control (when is it safe to cross the street, which direction should the robot go, clicking on a map).
- Monitoring: The human sees the sensor data and alerts or warnings sent by the robot (think security cameras). One human may be monitoring multiple robots.

The robot may or may not have audio or camera capabilities - the ability to record and speak to nearby pedestrians.

Even in a perfect tele-operation situation it is difficult to replicate the situational awareness a human has - everything from judging distances to hearing approaching cars. In particular, most cameras have a far more limited field of view and no spatial sound (only one microphone). On a noisy street it may not be possible for a human operator to clearly hear what is said to the robot.

Though some states have permitted fully autonomous robots, we suggest that some form of required monitoring is advisable, especially in the early stages of robot deployment in public spaces. Most robots should also probably work in shared-autonomy mode - where the robot can initiate basic safety protocols, like stopping if it runs into something, while also enabling a human operator to provide guidance. Again, proscriptive regulations that deal with every possible scenario are not warranted (or even possible) at this stage, but there are some basic goals that might be considered in regulation:

- Regulate a maximum time between breakdown to removal from hazardous situation (eg, robot stranded in the middle of the road). This is a mix of detecting the robot's failure and redundant systems, which provide backup sensing and movement even in the event of a software or wheel failure.
- Ability for pedestrians to communicate to a human through the robot. This can be beneficial both for reporting a robot's condition to the company and for emergency help for pedestrians in the environment.
- Something that might be considered by law-makers would be a compensation mechanism for the labour that members of the public might have to provide to address situations like broken down robots (to acknowledge the impact on people in these situations, and to financially motivate companies to mitigate and quickly respond to breakdowns and other impediments)

Use of public infrastructure: Who pays for it, who maintains it, who has access to it, and who owns the data?

So far the policy discussion has largely focused on the interaction between legislation, litigation, and sidewalk robot companies/operators. As outlined in the Environment section, many potential problems may be better addressed at the public infrastructure level, whether this means changes to the physical environment (access lanes, widening sidewalks, tunnels, instrumenting streetlights, etc) or creating a shared data environment (congestion levels, external cameras, real-time robot locations, dedicated robot "terminals"). This requires a conversation between technologists, city planners, the community, and the companies as to what the potential benefits and costs are for potential changes.

Although sidewalk robots will both be used by, and affect, people in the local neighborhood they are operating in, the balance of benefit and risk may (likely) not be evenly distributed. Just as Waze and other navigation apps dramatically increased local traffic in some rural areas[131] as commuters were directed to paths through them, sidewalk robots can fill up pedestrian pathways with deliveries from shops to consumers *not* in the neighborhoods they are driving through. Robots could be deliberately routed through specific neighborhoods, avoiding others. Some regulations are already touching upon location-based regulations – directing robots to some locations and away from others.[132] Some neighborhoods (because of a lack of sidewalks, bicycle lanes, curb cuts, or sidewalk maintenance) may not have access to sidewalk delivery at all.

Anyone familiar with efforts to increase cycling and pedestrian infrastructure in North American urban centers might also be imagining the types of practical challenges advocates of various other (non-robot, human-centered) uses of public space and infrastructure might face, not to mention the ways different interests might intersect with discussion about sidewalk robot infrastructural changes (e.g. cycling advocates might be concerned about sidewalk robots taking over hard-won bike lanes, but pedestrian advocates might see benefits in pushing forward more discussions about widening sidewalks etc)

---

[131] https://www.autoevolution.com/news/a-simple-sign-resolves-the-nightmare-waze-has-created-in-a-small-village-161123.html https://www.nytimes.com/2017/12/24/nyregion/traffic-apps-gps-neighborhoods.html

[132] E.g. Madison, Wisconsin (https://library.municode.com/wi/madison/codes/code_of_ordinances?nodeId=COORMAWIVOIICH11--19_CH12VECO_BIPLVE_12.753REPEDEDE)

There are many possible methods for cities to negotiate (or legislate) with sidewalk delivery companies to cover costs and changes to public infrastructure. We discuss below how technology might impact those decisions, again from a cost-benefit-privacy analysis. And as before, we will leave out discussing deliberate security hacks.

Structural changes

Stoplights, stop signs, and dedicated robot crossing lanes: Many stop lights already have some mechanism for police, fire, or rescue vehicles to trigger them, or the pressure plates used to recognize cars (and, sometimes, bicycles) waiting at the light. Some busy streets have a dedicated pedestrian crossing, where lights are triggered to enable pedestrians to cross. Some crosswalks have flashing lights built into the ground to notify cars of people crossing the road. There are two, relatively simple changes to these mechanisms that could increase safety both for the robots and drivers. Neither of these have obvious privacy downsides.

- Broadcast a signal indicating state (on/off, which direction). This completely by-passes the need for the robot to visually determine if the light, eg, is red. Note that this has potential benefit for visually impaired people as well, since this signal could be picked up by a phone[133]. This has a potentially important impact on safety, for minimal cost.
- Enable the robot to trigger the stop/crossing lights remotely - in essence, push the button. This would require additional security, licensing, and policy both to ensure that robots did not adversely affect traffic flow and to limit access to legitimate companies. This has less of an impact on safety, but would decrease wait times for robots (and may be necessary on street corners where the lights do not change without a button press).

Widening sidewalks, dedicated lanes or paths, curb cuts, painting sidewalk edges/cross-walks, implanting markers, trimming bushes, sidewalk maintenance

These are all non-technical methods for providing robots with more space to maneuver in, stretches where they have clear markings to delineate where it is safe(r) to move, or for improving localization. Robots traveling on clear, straight, well-defined paths can help

---

[133] What technology - wireless, blue tooth, radio - will impact how easily adoptable this is, since phones would need an adapter to pick up on the signal.

people predict where they are going and maneuver around them. Larger investments could include underground tunnels or overpasses to by-pass pedestrian heavy areas.

Widening sidewalks, trimming bushes, and dedicated lanes all come at some cost, of course, in terms of land use for driving and parking. While policy debates about who pays for such investments are complex and beyond the scope of a detailed analysis here, it seems worth noting that most sidewalk delivery robot companies in operation now are commercial in nature, meaning that it is private companies that benefit from the increasing use of these technologies. Where public investment is made into accommodating wider use of sidewalk robots, there should be substantial enough public benefit from these devices to justify expenditure, since the direct profit leaves the public sphere. Tax policy and congestion pricing could help balance the cost-benefit analysis of infrastructural investments. Investing in infrastructure is one important consideration for boosting the safety and potentially reducing the negative impacts of robots operating through public spaces, but it comes at a significant cost.

Conclusion - Reflections on Methodology - how do law-makers even approach articulating the trade-offs and concerns?

What is the process of discovering the questions we *should* be addressing when writing policy around technology? How should stakeholders work together to come up with better solutions? What should be legislated or regulated? What is the cost of those regulations? What other, broader issues (accessibility, bias, etc) will push the cost of *delayed* regulation on to someone else? What should be left to litigation versus legislation?

Throughout this paper we have identified a number of tensions and trade-offs between sidewalk robot regulation and technology. These include:

- Tension between simplicity of the device (or requirements for the device) and effectiveness – greater effectiveness often = more complexity
- Tension between data collection/privacy and physical safety – for robots to operate more safely they need to collect more refined data (to a point, there are limits on what is necessary to be effective), which more deeply engages privacy

- Tension between proactive and reactive laws – some issues require proactive consideration from law-makers; in other cases too much pro-active law might actually be counter-productive or problematic
- Robot gentrification – tension in where robots operate, and what benefit or cost this brings to people who live and make use of these spaces (or who make use of these devices, but don't have to experience their physical/environmental impact)
- In some cases – questions arise as to whether we even need these robots at all, given the costs/challenges? Why does delivery need to be automated? Who benefits and who pays? (e.g. Why not invest in better cycling and public transit infrastructure to facilitate delivery by people?)
- While not directly discussed in the above, a final tension relates to the use of public space – to what extent will laws permitting sidewalk robots normalize the use of public space by robots; or normalize prioritization of quick delivery over other possible uses of public space and infrastructure? Could increasing presence and popularity of robot delivery mean re-opening debates about aerial drone delivery to address some of the practical challenges of navigating shared spaces?[134] *Etc*

There are (broadly) at least five stakeholders that need to come to the table when discussing sidewalk robot regulation:

- Robot technicians/technical experts: What is feasible, from a technology standpoint
- Robot company business financial/marketing: Business models, cost effectiveness, marketing (as far as we can tell, these stakeholders have had a strong voice in law-making in many states so far)
- City planners: Knowledge of city infrastructure, development plans, possible control over zoning and congestion considerations
- Residents: Trade-offs between usability and public space uses, privacy concerns; including special interest groups/advocates
- City businesses: Potential users of technology, impacted by public space use.

One thing we learned in our own experience writing this paper is that one of the most effective parts of our methodology came from relationship and trust. We know each other, we could speak openly about our respective disciplines, brainstorm ideas, build on one another's tangent's. To do this effectively on a community/municipality level would require at least time and effort to establish networks of important stakeholders. Some of this work is

---

[134] Thanks to Woody Hartzog and Evan Selinger for a tought-provoking conversation about the potential legal/policy trajectory from sidewalk robot conjestion to drone delivery.

already done at the community level, but how to ensure the existing knowledge and expertise gets recognized in law making?

Overall, we suggest there is still much nuancing to be done in the laws that currently regulate sidewalk robots. It is our hope that deeper considerations will take place in the process of amending existing laws and adopting new laws across the U.S. and eventually in Canada.