

Micro Utility Devices - Observations from Transport Canada's Winter 2022 Urban Trial

July 2022

Transport Canada – Innovation Centre

Rielly Young & Mark Robbins¹

¹ A special thanks to the contributors to the project whose support allowed for this study to be produced Bryony McAllister, Sasan Ebrahimi, Tim Litchi, Jonah Robbins and the Innovative Solutions Canada program.

Micro Utility Devices (MUDs) or sidewalk robotics are an increasingly pervasive feature of city sidewalks world-wide. These devices are used to perform a wide variety of tasks at rates that are considerably more economic and more environmentally sustainable than existing technologies, though the emerging mode is not without concerns. As a new and rapidly emerging technological domain, MUDs operate within unclear and under-developed policy frameworks, leading to questions about safe and sustainable operations at scale.

As part of its ongoing work on emerging technologies in the transportation sector, the Transport Canada Innovation Centre has undertaken research to validate the current state of play in the sector as a first step towards policy development and federal standards which can be employed towards sidewalk robots. This work and research paper closely follows recommendations from a 2021 study commissioned by the Innovation Centre (Brail & Donald, 2021) which concluded that the Government of Canada should: 1) support pilot programs for sidewalk robots, 2) conduct policy labs, and 3) should actively socialize and disseminate knowledge about best practices in MUDs.

All three recommendations have been implemented by early 2022², with this report offering insights into our findings and directly reporting on observations of a MUD trial commissioned by the Innovation Centre with Swap Robotics through the Innovative Solutions Canada program over the course of Winter 2022, a subject covered in greater detail in the case study offered in Section 3 of this report. These overarching activities have also included policy labs, TC policy staff directly conducting research into standards and policies, directly supporting pilots with through novel procurement vehicles, and offering technical support, validation, and feedback to these pilot programs from TC's engineering teams.

This paper begins with an overview of MUDs and their operations, followed by an appraisal of the legal, policy and safety concerns associated with this emerging technology. The report then provides an overview of the direct observations and quantitative results of the Innovation Centre's Winter 2022 MUD trial in downtown Ottawa and ends with some conclusions about where future action for government might be appropriate and justified given the current state of play.

² In addition to convening stakeholders from across federal, provincial, municipal governments, as well as industry and civil society, our work has included a close collaboration with the Institute on Governance (an Ottawa think-tank) to conduct a multistakeholder policy lab. The results of this are available at (Miller, et al., 2022).

1. Introduction and Technical Overview

At the most basic level, MUDs can be described as remotely piloted robotic devices operating at surface level. These devices vary in size and operational scope, having the ability to deliver cargo, perform inspections, conduct vocational tasks including minor repairs and the like. To function and navigate their environment, MUDs use a combination of sensors – LiDAR, Accelerometers, Radars, Microphones, communication with nearby infrastructure (V2X) etc. – and cameras – visual spectrum, infra-red, computer-vision – to allow them to be remotely controlled, operate automatically with some degree of remote guidance in complex scenarios, or in some cases, operate altogether autonomously.

In the case of remote operations, the MUD is piloted by an individual that uses the device's cameras and sensors to allow for optimal situational awareness and blends the device's intrinsic ability to function autonomously in certain scenarios with the human operators' enhanced judgement to respond to challenging scenarios and edge cases for which the MUD was not pre-programmed. MUDs are typically equipped with high-powered CPUs capable of computationally complex tasks associated with artificial intelligence (AI). Depending on the level of autonomy, these devices may be equipped with communications technology to allow for remote operations and/or a hybrid onboard management system to execute decisions while overseen by a human monitor (Miller, et al., 2022).

While current trials will commonly see one human operator overseeing a handful of robots at any one time, depending on the task and domain, a remote operator can theoretically control dozens or hundreds of MUDs at one time (Basecelli, 2018). This makes MUDs an economically attractive alternative to crewed machines which must be staffed at a 1:1 ratio and must be designed around the ability to carry both an operational payload, and a human payload, the operator. By removing the onboard human operator from the design, MUDs are able to fundamentally rethink the design of small vehicles and their approach to payload. MUDs also typically operate using electric batteries and therefore a sustainably smaller environmental footprint compared to many existing means of surface transportation (Jaller, et al., 2019; Corno & Savaresi, 2020). MUDs also offer the advantage of being much smaller, more compact and fundamentally non-rival with existing road uses in congested urban centres, given that they operate on sidewalks instead of roads.

The most commonly raised commercial use case for MUDs pertain to last-mile delivery. Last-mile delivery is frequently identified as the single largest expense for logistical services businesses (Basecelli, 2018; Corno & Savaresi, 2020; Bakach, Campbell, & Ehmke, 2022), accounting for more than half of all

shipping costs, regardless of destination. Thus, the application of MUDs for last-mile delivery is an attractive potential solution to generate economic gains and find operational efficiencies. That being said, emerging use-cases are focusing on an increasingly wide variety of applications and scenarios such as plowing snow, persistent situational awareness, and infrastructure inspections.

Business development needs and techno-optimism often make pre-mature declarations of MUDs as being “autonomous” when in most cases the technical reality is far from this at the time of writing. In reality, there are very few autonomous MUDs in existence- the vast majority have some degree of remote piloting or remote monitoring. These devices are able to pilot themselves independently for short stretches of time but require persistent human oversight and intervention at regular intervals when the devices encounter circumstances which challenge their ability to respond safely and effectively in the absence of human intervention. Many subject matter experts prefer to use the term “automatic” since this suggests an ability to perform some tasks independently, while acknowledging that the system cannot operate in the absence of human oversight and intervention.

The Society of Automotive Engineers (SAE International) has developed a scale for assessing autonomous motor vehicle operations which adds some important nuance to the distinction between degrees of autonomy and automation and are directly applicable to MUDs. The SAE framework (table 2) is divided into six sections ranging from 0 to 5, where *Level 0* is fully manual operations and *Level 5* is fully autonomous operations. These levels of autonomy help to guide roboticists and policymakers alike to gauge the readiness of MUDs and ensure that the appropriate safety mitigations are put in place.

At a basic level, some of these safety mitigations are intrinsic and seemingly obvious. For example, a MUD operating at Level 2 autonomy- will by definition require operators to be onsite and in close enough proximity to have the option to easily intervene in the MUD’s operations. In this spirit, the safety of MUD operations should be gauged against the level of autonomy and the appropriate mitigations that have been taken for that level. For example, a MUD which operates at Level 1 autonomy can be deemed safe if its operations can safely and reliably be conducted within visual line of site of a human safety officer and with the regular intervention of an onsite operator. It would be inappropriate for example, to judge a MUD as a failure solely because it operates with the safety protocols associated with Level 1 instead of being able to operate in the complete absence of onsite operators, a feature associated with Level 5.

<i>Levels of Autonomy for MUDs</i>		
Level	Description	Time between intervention
0	Full manual operation	N/A
1	MUD operated withing visual light of sight	< 5 minutes
2	MUD operated with operator nearby or onsite	< 60 minutes
3	Several MUDs monitored by one on-site/remote monitor	8 hours
4	MUDs operate without on-site/remote monitor	3 days
5	Full autonomous operations – MUDs adapt and improve operational execution through ML	Extended operations > 3 days

While the above table demonstrates the potential range of MUD autonomy, in practice, there are few commercial operations that take place at Level 5. SWAP robotics' MUDs operate between level 2 and level 3 autonomy and are assessed accordingly in section 3.

As evidenced by the above table, the level of autonomy is gauged according to the amount of time between human interventions. The guiding principle for understanding operational complexity of MUDs should be based off relative level of autonomy. For example, if level 3 and up require appropriate signage to be displayed (as to inform the general public about MUD operations in area), such conditions would not be required for Levels 0, 1, & 2. The lower levels of autonomy require the presence of onsite/nearby visual observer. Thus, when a visual observer is not present, adequate signage shall serve as a proxy. As will be discussed further, trust and safety are intrinsically linked; thus, situating operational safety at the heart of MUD guidance shall serve to facilitate societal acceptance of this nascent technology.

2. Literature Review

The visibility of a MUD to other road and sidewalk users is a cornerstone of safe operations. For MUDs operating with a visual observer, the question of visibility is of less concern since the ability of a nearby operator to intervene can help mitigate the risks of operation. The proximity of the visual observer to the MUD is designed allow for quick intervention in the event of an incident. Where concerns arise in relation to visibility is around autonomous operation. As mentioned in the technical review, unlike typical vehicles which are large, have a height in the regular field of vision and operate with loud internal combustion engines, MUDs operate using battery power which can render them silent and have no fundamental design requirement that would guarantee they catch one's field of vision.



Figure 1 - SWAP Robotics MUD



Figure 2 - TinyMile MUD next to jogger

On one hand, from a design standpoint these features present a comparative advantage for MUDs over existing solutions since they are quiet and less obviously disruptive, allowing them to integrate more seamlessly into the urban fabric and contribute to a more livable city. On the other hand, this feature of MUDs can also present a potential safety concern because it makes it more challenging to difficult to notice these devices and react to them appropriately while operating in public settings, such as sidewalks. While there is a clear understanding that MUDs operating in public spaces need to be visible, there is a less clear understanding of what exactly that means from a design specification standpoint and how this can be done in a way that improves safety without detracting from the intrinsic benefits of MUDs being quiet and less disruptive than alternatives. A common approach has been to place a brightly colored flag or flashing light atop the MUD at eye level of pedestrians to catch their attention (See Figure 2).

While adding visual cues at eye-level offers some improvement, this feature alone is insufficient since they serve as a useful feature for able-bodied people with sight and offer little for those with visual impairment, cognitive disability or those who may be aware of the device but lack the physical ability to respond with dexterity. For individuals with visual impairments, an audible indication of an MUD could serve as an appropriate safety feature however, this too is limited in its effectiveness since new audio cues can be disorienting for the blind and notably, offering no improvement in perception for the deaf-blind. Some of the solutions proposing to fill this gap are technological, such as open-source transponders that can interact with phone-based accessibility apps. Other mitigations can take the form of appropriate governance and awareness mitigations, such as restricting access to area where clear signage can be displayed, geofencing or limiting the time of permitted operations to moments where pedestrian traffic is infrequent.

The latter mitigation is one of the more obvious and easier to deploy in the short term since the sites where MUDs are being deployed will greatly inform the risk profile of the operation and safety considerations that will need to be adapted. Whether a site is uncontrolled, controlled or semi-controlled will help to determine what safety mitigations can take place and the sufficiency of these measures to bridge the divide between the navigation stack of the MUD and the profile of users. For example, a controlled setting such as inside an airport terminal may be able to design the physical space in a way that aligns with the capability of the technology, such as the Calgary Airport which has segregated and clearly demarked zones where semi-autonomous devices are able to navigate with a marginal risk of directly interacting with pedestrians.

If an MUD were by contrast operating in a quasi-private semi-controlled environment with fixed access points such a university campus or shopping centre, there would be realistic opportunities to disseminate information on MUD operations to site users at the point of entry, raising awareness and perceptibility of others using the space. Such areas could also offer supports to site users at the point of entry to the site (guides, chaperones, etc.) to ensure that those with reduced mobility or audio-sensory perception are still able to navigate in the semi-controlled environment comfortably and safely.

Fully uncontrolled environments, such as the public sidewalks in the downtown core of a major city, offer the least ability to implement these sorts of safety mitigations since there are many access points and few if any restrictions on who is able to enter and exit these areas. In these situations, the best governance-oriented mitigation may be geofencing the MUDs themselves. At its most simple, this creates zones where MUDs can operate and others where they cannot based on an assessment of the

risk associated with these zones. More nuanced approaches could see certain sizes of MUDs permitted in certain areas but not others, or access permissions changing by the time of day and the like (i.e. no MUDs allowed near subway entrances at rush hour).

MUDs in the uncontrolled public domain are most exposed to legal liability as members of the public may have little to no knowledge of their operations prior to first contact. For instance, if a delivery MUD was operating in a neighborhood, any number of interactions can occur that cannot necessarily be mitigated by the operators. Therefore, the visibility of the MUD, the intrinsic intelligence of the device, and the integration of human observers (onsite or remotely) become more important in uncontrolled public sphere of operations. In the case of an MUD in a quasi-public sphere of operations, there is less of a concern with legal liability since information can be disseminated and mitigations more readily undertaken. This would allow individuals to exercise proper caution when entering the MUDs operational environment. Thus, the same MUD operating in a controlled environment has a lower risk profile than one operating in an uncontrolled environment, because the potential for an individual not associated with the MUDs operations to encounter the device is inherently lower than that of an uncontrolled environment. This approach mirrors some of the well-established principles at play in regulation of airspace and integration of drone operations into different risk environments.

Hoffman & Prause (2018) distinguish between two types of liability MUDs would face while in an uncontrolled environment. The first is liability for torts inflicted by traffic incidents. This has been a common requirement for municipal/state pilot projects pertaining to MUDs in the United States. Washington DC's regulation of *Personal Delivery Device Act* (PDDA) required applicants to procure at least \$250,000 of liability insurance to be permitted to operate within the district (Basecelli, 2018). The Washington DC model has been subsequently adopted by several other states, including Florida, Idaho, Virginia, and Washington State. The second type of liability that MUDs would encounter pertains to product liability – the legal liability manufactures/retailers face when their product causes harm to a member of the public.

While these liability issues are important to be aware of and for manufacturers to respond to, it is unlikely that liability issues will prevent the continued progression and adoption of MUDs as a technology. Though MUDs may well represent a novel application for which there is little specific policy or regulation at the time of writing, there are indeed many well-established legal processes and judicial interpretations which can easily be read-in to scenarios applying to MUDs. Take the example of a MUD operating on municipal sidewalks for the purpose of snow removal. If the MUD causes damage to

persons or property due to negligence, the harmed individual can seek legal recourse with the municipality, the manufacturer, or both. The snow removal example is extremely prescient with the backdrop of a 2021 Supreme Court decision that found snow removal to be a “core policy decision” for municipalities and are therefore not immune from negligence claims for unplowed snow (see *Marchi vs. City of Nelson*, 2021). For municipalities, robotic maintenance like snow plowing touches on two types of legal liability for which they are responsible, that of introducing new products on one hand and that of the residual liability that comes with leaving snow unplowed. Indeed, interviews from this research suggest that municipalities commonly respond to snow events in a timely manner in a way that covers a small fraction of the areas where a fall due to lack of plowing represents a legal liability for the municipal government. Similar observations can be noted for other maintenance issues which represent a residual liability, such as crack inspection and repair for sidewalks which present a tripping hazard. This is to say that the status quo can be understood to also present a risk.

A common, but less widely known concern pertaining to MUDs relates to the collection of data. As noted by Hoffmann & Prause (2018), MUDs require the constant collection of data from a range of sensors and cameras. Though the privacy question is frequently raised in academic discourse, this is in many ways a “solved problem” since there are existing regulations and codes of practice pertaining to data collection and privacy which have a long history of being read-in to new technology areas. MUDs are subject to existing data privacy laws and should be operated according to the best practices of privacy by design, such as minimum collection of data, differential privacy, federated learning and the right to be forgotten. While more collection may take place as MUDs become more pervasive, this is similarly true as streetlights are gradually replaced with smart infrastructure, or as the stock of motor vehicles on the road becomes increasingly “smart”. Though MUDs may attract more attention, it seems unlikely that their increasing pervasiveness will require a tailor made or entirely novel solution more than diligence in compliance with existing regulatory standards and policy prescriptions for privacy.

	OECD Privacy Principles
Principle 1	Collection Limitation
Principle 2	Data Quality
Principle 3	Purpose Specification
Principle 4	Use Limitation
Principle 5	Security Safeguards
Principle 6	Openness
Principle 7	Individual Participation
Principle 8	Accountability

Nevertheless, some authors identify privacy as a major concern for the introduction of MUDs in public spaces (Mintrom, et al., 2021). Depending on the level of autonomy, MUDs require a plethora of data to enable situational awareness and operation safety. As the level of autonomy increase, the greater the reliance on data

collection is for MUD operations. This mass collection of data can lead to ethical concerns regarding its collection, storage, and usage. One suggestion to mitigate public concerns in respect to privacy is to proactively broadcast not only what the device is doing, but also what information is being collected. This approach was piloted in the short-lived “Sidewalk Labs” Smart City district in Toronto, where a system of labels and signage was developed to signal what data was being collected in each district and how it might be processed or shared. These collection protocols could similarly be geofenced and timed as required by the safety, security or privacy considerations of the locations where they operate.

Other jurisdictions have taken the approach of implementing guidelines for building trust between individuals and private firms operating devices that collect data as part of their regular operations. Examples of these frameworks include the EU Commission’s Guidelines for Trustworthy AI and the OECD’s Privacy Principles (Table Three) A study of European Union robotics companies (n=10) found that nearly all of the privacy principles listed in Table 2 could be found in the firm’s privacy policies [P1 through P7] (Chatzimichali, et al., 2020). Indeed, social acceptance of MUDs is an important pre-requisite for commercial adoption of the technology to occur and the ability of governments and industry players to earn public trust in the operation of MUDs will determine the success of the sector in Canada. Research conducted by UPS suggests that customers are not only incredibly open to the idea of MUDs for last mile delivery, but they would be willing to pay extra for MUD delivery if they believe their package will arrive sooner (Bakach, Campbell, & Ehmke, 2022). This suggests that the public’s willingness to accept MUDs is correlated with not only the service being conducted but the expected net benefit to the individual, which is generally understood to be a contribution to net prosperity.

There simultaneously some concerns about how these devices will impact equity. Jaller, et al., (2019) argues that concerns from labor groups often express fear and contempt towards MUDs with the perception that these devices will result in a reduction in the workforce. Such hesitation was demonstrated by a UPS study of their workforce; when the firm attempted to implement a pilot program that used drones to perform certain delivery tasks, they were met with significant resistance from Teamsters Packing Division (Jaller, et al., 2019). On the other hand, many of the industries operating in the transportation, logistics and maintenance domains where MUDs and growing in prevalence are similarly faced with severe labour shortages reducing their ability to operate effectively. This signals a need for constructive dialogue about precisely what functions will be subject to replacement and how the benefits of increased automation will be shared between industry, labour and users.

3. Case Study: MUD Sidewalk Snowplows at Tunney’s Pasture (Swap Robotics)

Through the Innovative Solutions Canada – Testing Stream, Transport Canada piloted a suite of MUDs from *Swap Robotics in downtown Ottawa. This took place* on federal lands at Tunney’s Pasture with the strategic objectives of improving TC’s technical understanding in this domain, building capacity at lower orders of government and advancing towards the development of federal standards for future trials. At a tactical level, the goal of the pilot was to evaluate the technology as a forward-looking solution for snow removal on federal lands and to validate the safety of autonomous/semi-autonomous functionality.

To confirm safe operations of the SWAP Sidewalk Snowplow, two types of tests were conducted: base cases and edge cases. Base cases can be described as typical interactions that the robot can expect to encounter at any moment during snow removal and for which the MUD has been designed to respond. Edge case testing sought to understand how the MUD would react to interactions with environmental/human factors that are likely to occur infrequently, and thus may not have been addressed as part of the core design features of the device. The purpose of edge case testing is to probe for failures to help inform future governance mitigations that can be undertaken by governments, and technical developments that can be undertaken by roboticists and other manufacturers.

Base Cases:

#	Test	Measurement	KPI (Key Performance Indicators)	Result
1	Confirm that the MUD is safe enough for active commercial snow removal operations through demonstration of safe grid crossings while conducting commercial operations.	Robots demonstrate safe street crossings that are as safe, or safer, than traditional equipment street crossings.	Confirmation of the test requires 18 out of 20 street crossings to be done so without safety concerns arising.	Successful completion of 20/20 street crossings
2	Confirm that the MUD is safe enough for active commercial snow removal operations through demonstration of awareness of, and safe interaction, with pedestrians and animals.	MUD demonstrates awareness of, and safe interaction with, pedestrians and animals commiserate with safe commercial snow removal.	Confirmation of the test requires the MUD to demonstrate awareness of pedestrian $\geq 80\%$ of the time (ie., the MUD slows/stops when encountering pedestrian/animals).	Successful identification and response in 17/20 standard pedestrian and animal scenarios
3	Confirm that other safety-related features work as intended and are operational.	Visual observation of all safety features during operations: [1] all emergency-stop	Confirmation of the test requires all listed safety features during operation.	Successful confirmation of all safety features

		buttons, [2] glowing lights around emergency-stop, [3] Reverse lights, [4] reverse beeper when reversing, [5] Signal lights when turning, and [6] brake lights.		throughout 2 months of repeated testing.
4	Determine if the MUDs are ready for longer term, multi-year sidewalk snow clearing & salting contracts with Government of Canada departments from a cost, output & productivity perspective.	Compared against existing means of federal snow removal.	Confirmation of this would require SWAP's MUD to be more cost-effective than traditional means of snow removal.	Positive assessments from Government of Canada users across multiple departments.
5	(5) Compare the benefits and drawbacks of electric powered snow removal equipment to gasoline equipment.	A graded rubric of several areas of equipment performance compared to two pieces of legacy equipment, including: [1] noise, [2] run-time, [3] maintenance, [4] environmental impact, and [5] total cost.	Qualitative assessment of the benefits and drawbacks of electric powered snow removal equipment compared to traditional methods of snow removal (gasoline powered equipment).	Compared to traditional methods, electric powered equipment: [1] is quieter; [2] is longer lasting; [3] has lower maintenance cost; [4] has a negligible environmental impact; and [5] is approximately half the financial cost.

All base cases were tested at least two separate times – some cases were tested more than two times. At no point during the testing period did the MUDs operate without physical guardians walking alongside the device. These guardians monitor the safe operations of the MUD. In the event of a potentially dangerous situation, the on-site guardians can intervene and cease operations immediately. For instance, confirmation of all safety features was continuously monitored throughout the testing period. At no time during the testing period did the MUDs safety features fail to operate. As such, most of the testing of base cases was done so at level-2 autonomy.

The first base case tested sought to determine whether the MUD could successfully complete a street crossing. Street crossings present the highest risk profile in terms of basic operations. Not only does the MUD have to compete for space with vehicles, but also other pedestrians. Pedestrians on sidewalks are

typically in a constant state of motion and thus only take up space on any given part of the sidewalk for a brief moment. However, people waiting to cross the street often congregate at crosswalks; thus, making it more difficult to operate around for MUDs. Accordingly, the testing parameters for this base case were high: successful operations required eighteen out of twenty street crossings to be done so without safety concerns arising. This test was successfully completed twenty out of twenty times.

The second base case sought to determine whether the MUD could operate in a commercial environment while safely interact with pedestrians and animals. Situational awareness is imperative to enable saving cohabitation between pedestrians, animals, and MUDs. The MUD was able to independently demonstrate situational awareness in all standard pedestrian and animal scenarios in a variety of winter circumstances and against different background settings, angles, and silhouettes to ensure robustness. Following the success in this base case, additional scenarios were tested in the seventh edge case.

The third base case sought to confirm the operational functionality of the MUDs safety features: [1] all emergency-stop buttons, [2] glowing lights around emergency-stop, [3] Reverse lights, [4] reverse beeper when reversing, [5] Signal lights when turning, and [6] brake lights. All safety features were tested and operated without incident on many occasions throughout nearly 2 months of active testing.

The fourth base case is a qualitative assessment of whether these MUDs are suitable for integration into regular federal snow plowing operations. Feedback from existing snowplow service operations, site managers and other government departments consulted were overwhelmingly positive. Snowplow service operators specifically found that the MUD plows were a useful tool in their present state that could be used augment existing services and replace some manual plowing. The technology continues to be improved upon with future generations of these MUDs having greater ability to operate independently, something which was felt to improve their usefulness in any integration scenario.

The fifth and final base case sought to evaluate the SWAP MUD against two standard snow-removal systems. As demonstrated in the below table, the SWAP MUD was evaluated against traditional methods of commercial snow removal with comparable functionality. Labour was calculated at \$35/hr for 300 hours; therefore, average labour costs \$10,500. Average maintenance costs for traditional snow removal methods range from \$3,000-\$7,000.

	SWAP MUD	Snowrator	HOLDER Sidewalk Plow
Noise	70 dBA	95 dBA	120 dBA

Run-Time	4 hours between charges; batteries can be swapped out in under 5 minutes.	Honda GX390 uses 92 gal/hour @ 3600 rpm with an approximate run time of 1 hour 40 minutes	Kubota 4-cylinder 4-stroke turbo diesel uses 1/2 gal/hour @ 2400 rpm with an approximate run time of 30 hours (tank holds 17 gal)
Maintenance	Approx. maintenance ratio 1:30 (1 hour of maintenance per 30 hours of operation)	Approx. maintenance ratio 1:10 (1 hour of maintenance per 10 hours of operation)	Approx. maintenance ratio 3:10 (3 hours of maintenance per 10 hours of operations)
Environmental cost	Null when charged on a clean energy grid	Honda GX 390 internal combustion engine	Kubota 4-cylinder 4-stroke turbo diesel
Total Cost	Average seasonal cost: \$15,000	Unit price: \$14,000 Total cost: \$28,500 ³	Unit price: \$74,000 Total Cost: \$94,500 ⁴



Figure 4 - Boss SNOWRATOR



Figure 3 – HOLDER Sidewalk Plow

³ Unit price + labour + maintenance + fuel = \$14,000+\$10,500+\$3,000+\$1,000 = \$28,500

⁴ Unit price + labour + maintenance + fuel = \$74,000+\$10,500+\$7,000+\$3,000 = \$94,500

Edge Cases:

Edge cases were developed to push past the proven limits and claims of the existing technology stack to identify weak points that MUDs are likely to encounter when operating in the public domain and at higher levels of autonomy. The edge cases yielded mix results, with Swap successfully identifying and responding to certain edge cases with relatively high levels of reliability, while other edge cases were not completed successfully. To be clear however, in all these cases the device was being asked to stretch beyond the stated capabilities of the manufacturer in order to identify weaknesses in the device after safe operation at level 2 autonomy had been successfully demonstrated. In spite of the shortcomings identified in some of the below edge cases, this MUD was nonetheless assessed as being able to operate reliably at level 2 autonomy.

#	Test	Measurement	KPI (Key Performance Indicators)	Result
1	Confirm that the MUD is safe enough for active commercial snow removal operations while in conflict with a pedestrian.	Robots demonstrate safe interactions with pedestrians while in conflict.	Robot ceases operations with 5 feet of pedestrian without intervention from guardian.	Was successful 75% of the time, with guardian intervention otherwise required.
2	(2) Confirm that the MUD is safe enough for active commercial snow removal operations while in conflict with a wheelchair.	Robots demonstrate safe interactions with pedestrian in wheelchair while in conflict.	Robot should cease operations within 5 feet; defer to WC without intervention from guardian.	Was successful 75% of the time, with guardian intervention otherwise required.
3	(3) Confirm that the MUD is safe enough for active commercial snow removal operations while in conflict with a wheelchair while approaching on a curve/blind turn.	Robots demonstrate safe interactions with pedestrian in wheelchair while in conflict on curved track (sidewalk).	Robot should cease operations within 5 feet; defer to WC without intervention from guardian.	Successful in both cases (rightward turn & leftward turn) where this was tested.
4	(4) Confirm that the MUD is safe enough for active commercial snow removal operations while in conflict with a wheelchair while approaching a crosswalk.	Robots demonstrate safe interactions with pedestrian in wheelchair while in conflict while approaching a crosswalk.	Robot should cease operations within 5 feet; defer to WC without intervention from guardian.	Successful in the one case where this was tested.
5	(5) Confirm that the MUD is safe enough for active commercial snow removal operations	MUD demonstrate safe interactions with small pedestrian (toddler) while in conflict.	Robot ceases operations with 5 feet of pedestrian without intervention from guardian.	Unsuccessful. MUD did not recognize toddler. Computer vision system only

	while encountering a toddler (or equivalent small human <20kg and < 1m in height.			recognizes adult-sized shapes.
6	(6) Confirm that the MUD is safe enough for active commercial snow removal operations while encountering a stroller.	MUD demonstrates safe operational functionality while encountering a stroller.	MUD ceases operations with 5 feet of stroller without intervention from guardian.	Successful 2 out of 5 times.
7	(7) Confirm that the MUD is safe enough for active commercial snow removal operations while encountering a dog.	MUD demonstrates safe operational functionality while encountering a dog.	MUD ceases operations with 5 feet of the dog without intervention from guardian.	Unsuccessful 2 out of 2 times.
8	(8) Confirm that the MUD is safe enough for active commercial snow removal operations while interrupted with a construction pylon – Forward.	MUD demonstrates an adequate level of situational awareness when encountering forwards-facing physical obstruction.	MUD ceases operations prior to contact with physical obstruction without intervention from guardian OR ceases operations after physically encountering the object (under 2 seconds).	Activation of the forward-facing pressure sensor terminated operations within 2 seconds of physical contact
9	(9) Confirm that the MUD is safe enough for active commercial snow removal operations while interrupted with a construction pylon – Backwards.	MUD demonstrates an adequate level of situational awareness when encountering backwards-facing physical obstruction.	MUD ceases operations prior to contact with physical obstruction without intervention from guardian OR ceases operations after physically encountering the object (under 2 seconds).	Activation of the backwards-facing pressure sensor terminated operations within 2 seconds of physical contact

The first edge case sought to demonstrate the safety of the MUD while abruptly being confronted by a human. The MUD was successful on three out of four attempts at this edge case. In the three successful attempts, the device ceased operations when a person came within five feet of it. On the one unsuccessful attempt, the MUD recognized a potential obstacle and slowed down before it was within a five-foot range. However, once the device crossed the five-foot barrier without completely stopping, the remote guardian terminated operations and the device came to an immediate halt.

The second edge case sought to demonstrate the safety of the MUD while encountering a pedestrian in a wheelchair. This case was of particular interest, as individuals with mobility issues are one of the groups most at risk of harms caused by adverse interactions with MUDs and have a limited ability to

make way in the event of a MUD which is failing to identify the user. Similar to the previous test, the MUD was able to demonstrate safe interactions on three of four attempts. For the first test, the MUD failed to identify the person in a wheelchair and did not begin to slow down as they approached one another. Before the MUD was five feet away from the wheelchair, the remote guardians terminated the test. The test was completed three more times and successfully identified the person in a wheelchair and stopped before breaching the five-foot barrier. The test was also completed in a crosswalk setting. SWAPs Standard Operating Procedures dictate street crossings are controlled by the remote operator – meaning autonomous street crossings would not occur during real-world operations. This is consistent with other MUD companies such as Serve Robotics and TinyMile. Nevertheless, the ability for autonomous street crossing was still tested.

Similar to the second edge case, the third edge case sought to demonstrate the safety of the MUD while encountering a pedestrian in a wheelchair at a sharp corner/blind turn. This scenario was tested four times total – two on a rightward turn and two on a leftward turn. For the rightward turn, the first test was ultimately unsuccessful; while the MUD slowed down at an approximate distance of ten feet, once the device reached the five-foot barrier the test was terminated by the remote guardians. The second test was successful – the MUD reduced speeds within 5-10 feet of the wheelchair and ultimately stopped within 5 feet.

The fifth edge case sought to demonstrate the safety of the MUD while abruptly encountering a toddler standing alone in its path. The first test was terminated by the remote guardian after the device came within five feet of the simulated toddler. The second attempt was slightly more successful than the first; the MUD recognized the object in its path and slowed its speed. However, the test was ultimately unsuccessful as the MUD failed to stop, when it reached the five-foot barrier the test was terminated by the remote guardians. Ultimately neither edge case test was successful. It is not clear whether this failure was due to the size of the human silhouette or if the cladding in winter attire increased the difficulty in detecting the silhouette.

The sixth edge case sought to demonstrate the safety of the MUD while encountering a stroller in its path. This case was tested five times in total: twice with the stroller facing the MUD as it approached, twice with the stroller facing away from the MUD as it approached, and once with the stroller facing perpendicular to the MUD as it approached. The MUD was only able to stop without remote guardian intervention on two of the five tests. Both successful tests were with the stroller facing away from the MUD as the device approached it. The exact reasoning behind this testing discrepancy is not entirely

known, however it is assumed that the backwards facing was successful as a result of the tallest point on the stroller (ie., the handles) were on the rear of the stroller; these handles might not have been within the computer vision's line of sight prior to termination by the remote guardian.

The seventh edge case sought to demonstrate the safety of the MUD while encountering a dog. This case was tested twice and was not successful in either attempt. As the MUD approached the device failed to slow its speed and was ultimately terminated by the remote guardian.

The eighth and ninth edge cases were similar; in that they were trying to simulate the conditions of the previous test (dog) while confirming functionality of the front and rear bumpers. Each bumper is equipped with a pressure sensor meant to terminate operations if triggered. To simulate the previous test, a construction pylon was used. The pylon was approximately 3 feet in height and weighted less than 15 pounds. The functionality of the front and rear bumpers was tested twice each for a total of four tests. As expected, on all four tests, the MUD's computer vision did not identify the pylon and stop prior to contact. Nevertheless, both the eighth and ninth edge cases were successful, as the device halted operations within one second of the object contacting the bumper. The KPI for both edge cases were a stop-rate < 2 seconds; thus, the MUD passed both cases.

4. Lessons Learned & Recommendations

Sidewalk robotics are an increasingly important part of Canada's transportation system and will only grow in their significance. Close consultation with existing expertise, combined with TC's own real-world trial undertaken with Swap Robotics have shed light on some of the issues at play and hurdles that will need to be overcome as this sector continues to mature. This section is designed to offer insight into some of the lessons learned from this research and real-world trial, and to offer some recommendations for future action.

Recommendation #1 - Capacity-Building

Transport Canada should leverage its expertise and convene other experts in the ecosystem to help increase domain expertise among municipalities and provinces and build regulatory capacity in the ecosystem.

Though existing regulatory standards and research conducted by Brail and Donald (2021) both suggest that the federal government has little formal authorities to directly regulate MUDs operate, Transport Canada does have an outsized ability to impact the sector through its scientific and R&D leadership, as well as its power to convene disparate groups to build capacity across the system. While some larger provincial governments and municipalities may well have the ability to conduct their own R&D and experimental trials which are ultimately needed to inform dedicated regulations, a great many more simply lack the organizational size and scale for dedicated regulations to be realistic. When municipalities lack the bandwidth to thoughtfully engage with a nascent technology, there is always a risk of either inaction or, at the opposite extreme, the imposition of a hasty ban. A similar dynamic was observed with the initial roll-out of ride sharing services. In a best-case scenario, this asymmetric capacity and access to R&D capabilities results in a handful of municipalities and provincial governments being "rule makers" with the vast remainder becoming "rule takers", copying regulatory frameworks from other jurisdictions or from companies themselves.

Transport Canada can play a constructive role here by conducting applied R&D in the MUD space and making the results widely available as possible to all actors in the ecosystem. There is a clear value to federal R&D and experimentation being conducted in a way that maximizes openness since the results of this work is non-rival and can lend to the capacity and knowledge of any number of regulatory players within the Canadian ecosystem. Given the magnitude of the challenge coming from the

projected exponential increases in the number of MUDs in operation, it is incumbent on Transport Canada to continue to conduct R&D in this space, aggressively disseminate results and convene implicated players around best-practices and avoidable failures. In its efforts at wider capacity-building and research dissemination, TC should be actively inclusive of the MUD industry and related research organizations since they can be an important source of capacity and their expertise should be widely leveraged.

Though MUD-like technologies are nascent in the context of the public sphere of operations, the technology itself has ample precedence in private settings like factories and industrial venues. Many industry groups – such as SAE, ISO Tech Working Group, among others – have established SOPs for autonomous or semi-autonomous technologies. While regulators should not solely defer to industry, it is important to include the insights garnered by those with the most firsthand knowledge of the technologies. Thus, collaboration with industry to develop thoughtful guidance should be considered an asset when looking to develop safety standards and regulations. This includes issues of operational uniformity for design features such as the height of a flag, the volume of an operational indicator, and the color of a flashing light. Many of these issues are fundamentally “solved problems” though the information about these solutions is poorly socialized, particularly among regulatory stakeholders. Simple methods for promoting innovation and increasing public acceptance of the technology.

Recommendation #2 - Real-World Trials

More real-world trials are needed, not less, and all levels of government should strive to foster them and widely disseminate the results.

The current state of play suggests that more real-world trials will be necessary and are likely to be conducted with increasing frequency across the transportation ecosystem. There are few substitutes for real-world testing when it comes to emerging technologies in large part because these trials identify failures that are difficult to conceive of in simulated or lab settings. One clear example noted in recent field testing was the use of QR codes, which are sometimes placed prominently on MUDs to allow the public to gain a better understanding of its operational function. This practice aligns with the academic literature on technological acceptance and public trust; however, it also unintentionally has an outsized detraction from safety since it brings bystanders close to an operational MUD as they try to scan the QR

code. Observations like these from real-world deployments are key to informing future designs of MUDs so they are as safe as possible, as well as to inform site governance and operational standards deployed by provinces and municipalities.

It is crucial that these kinds of results from trials be made as broadly available as possible so that lessons and best practices can be used to widely inform municipalities and provinces, the vast majority of which lack the scale of operations necessary to conduct these trials themselves. This practice is beneficial to government authorities and innovators alike. For innovators, there is an efficient, streamlined process for technology and feature validation and input on how designs can be as safe and accessible as possible. For government authorities, these pilot projects offer the opportunity to confirm the safety of an innovation before wider commercialization and allow for current/new regulations/policy's to be evaluated in a quasi-real-world environment. This immense need for trials takes place against the backdrop of several major Canadian cities (including Toronto and Ottawa) introducing temporary bans on all uses of sidewalk robotics, including in trial formats. This is not conducive to innovation nor to the dialectic between regulator and innovation that it ultimately crucial to driving safe adoption of these technologies in the public domain. Such bans should be reconsidered.

A renewed emphasis on controlled real-world trials and can be supported by increased activity at TC, regulatory renewal at provincial and municipal levels of government, and the targeted creation of test beds for wider use. The targeted creation of test beds for real-world trials serves to benefit more than just capacity building: this is the opportunity to provide rules and governance of future pilot projects and deployment by other levels of government. With this being said, test beds and trials are an interim step and not an end state. The growing prevalence of test beds and real-world trials needs to take place as a step towards a system of governance that aligns hazard with the potential of harm and mitigation efforts. This will include risk matrices for certain zones and MUD classes (sizes, speeds etc.), acceptable mitigations to decrease the risk profile of a deployment and stated conditions under which a MUD may or may not be deployed. All of this information should be made widely accessible to the public in the interests of promoting compliance and building a viable social compact around MUDs.

Recommendation #3 - Experimentation with Geofencing

Comprehensive geofencing according to use and built environment should be considered as part of developing a living and adjustable system of rules to govern urban MUDs.

The most immediate use for geofencing is to create zones that define the acceptable area where trials may take place and to bound MUDs to those areas. To receive permission to operate in a jurisdiction, these robots must be functionally hard-wired to obey the rules for operation set out in the geofencing, or they deactivate. Geofencing sets aside specific areas in the urban fabric for specific purposes and prevents robotics (or other devices) from breaking the rules associated with the geographic area. This can be used to allow for innovators to test their devices under strict circumstances, effectively creating an artificial semi-controlled domain appropriately aligned with risk, mitigations, and even perhaps informed consent for those entering the geofenced area.

While geofencing is an important immediate step to permit the safe integration of real-world trials into the built environment, intelligent and comprehensive geofencing of jurisdiction could prove to be more than just an interim measure, but a steppingstone to an end state. Geofencing should be used not just to designate Green Zones and Red Zones (where operations are fully permitted and fully forbidden respectively) but can be used with greater nuance to align with times of non-disruptive operation, size of vehicle, speed limits and the like. For instance, some high traffic pedestrian areas may be closed to MUDS during rush hour, and open to the largest and fastest (and thus most disruptive) MUDs in low density areas between 2am and 4am. It is possible to add much greater flexibility to a body of rules centered around geofencing than the alternatives, geofencing can also be altered in real-time to create experiments or to adjust rules as new information or best practices come to light.

By using geofencing as the central site for applying regulation, governments can functionally put into practice a body of rules that are mandatory for the robots operating in a given space, greatly increasing compliance while also dramatically reducing the cost of enforcement. In turn, comprehensive geofencing can gradually be attuned to urban features with a micro level of precision and optimizing, incorporating details such as the size of sidewalks, the likely presence of pedestrians and VRUs, nearby built obstacles, times that schools or sports venues are likely to crowd sidewalks as so on. This hyper-optimization potential helps greatly to ensure that the robots are operating as safely as possible and in a way that is both readily adjustable to changing circumstances and accepting of new regulatory and technological innovations. Learnings from the operation of different devices can be easily incorporated

into practice by constantly iterating upon the geofencing; attuning it as lessons are learned and seeing an immediate adjustment in the robots' behavior. Early geofencing efforts to permit "safe zones" for trials should thus be viewed as a first step to be continuously built upon to expand the safe operational use zones for MUDs.

Recommendation #4 - Public Goods in Robotics

As an emerging technology the development of MUDs depends on several types of technologies groupings that are in rapid evolution including various kinds of sensors, computer vision, geospatial awareness, and the like. While a common trajectory for these technologies is for individual designs to be developed and commercialized, other sub technologies benefit from being more openly shared and distributed. A common example might be open-source software, where the development of common code that is available to everyone spurs greater innovation across the wider network. The deployment of open-source technologies goes far beyond the level of theory; Pixhawk open-source autopilot for example has become the de-facto standard for many classes of small drone. For the emergent field of MUDs, there remains an open-ended question about which subsystems and technologies will reach maturity as proprietary technologies, and which will develop as openly available, and perhaps even develop as a form of public good.

One proposal from Miller et al (2022) suggests the pooling of geospatial and locational data (images and mapping of sidewalks) could be a useful tool to facilitate the safe progression of the industry. High resolution and regularly refreshed urban geospatial data would dramatically improve the reliability of MUD navigation systems. This data is non-rival and has the best system level impact when shared as widely as possible since access to this information will improve safety and effectiveness at no additional cost per user. Geospatial data for urban navigation therefore can be viewed as a textbook public good. Machine learning algorithms – like the ones that enable the application of computer vision – for promoting safe navigation may well fit under the model of a public good as well. Detection of objects that may be encountered by a MUD requires a vast amount of training data, and these kinds of data sets typically focus on common objects to the exclusion of objects that are less common but no less important to avoid, such as wheelchair users or unaccompanied toddlers. Federal, provincial, territorial

and municipal governments need to convene around this subject and weigh in for the creation of public goods where appropriate.

Finally, governments play an important role in the development of standards that can be enacted and socialized to support an even playing field through certain minimum standards of safety and interoperability. Scholars agree that the goal of regulation for autonomous robotics should seek to adopt common standards for data collection and human-robot interaction that can apply to all market participants, rather than directly regulating autonomous robotics as a discrete and separate subject (Chatzimichali, et al., 2021). Though much has been done in this space by innovators to develop working standards, where these ad hoc practices do exist they are poorly socialized with regulators at all levels of government. We are very far away from a de facto industry-led standard.

Recommendation #5 – Accessibility by Design

Public trust and understanding are clear requirements for the widescale adoption of emerging technologies including MUDs. More needs to be done by innovators and governments alike to build trust and develop a viable social compact surrounding the use of MUDs in public spaces. This will partly be built by greater exposure of the public to MUDs, but it's also crucial that deployments of MUDs be underpinned by an ever-greater adaptation of MUDs to their environment, making them an ever more seamless part of the public domain. While MUDs may do well in interactions with run-of-the-mill pedestrians, fixed infrastructure and cars, a better sign of meaningful continuous improvement is the ability of MUDs to respond and adjust to edge cases where their performance is significantly less consistent, in particular for the recognition of, and response to, people with disabilities. Accessible design will continue to be crucial factor in driving public acceptance of MUDs since the net benefits of MUDs must be widely and fairly distributed in accordance with the increasing emphasis on universal design in Canada. The present offers a critical juncture in that the design and production cycle of these devices since main remain sufficiently nascent that they can easily incorporate novel technologies and best practices for accessibility, rather than being forced to commit to expensive or implausible retrofits.

Though there have been some concerns raised about the impact of MUDs on the accessibility of the transportation system, technology is value agnostic at its core and the impact of MUDs on accessibility will ultimately reflect the demands for accessible design that are placed on it by governments,

innovators and the citizenry. In other words, so long as innovators are aware and meaningfully responsive to the principles of accessibility by design, there is every reason to expect that MUDs will have a net positive impact on access in Canada. For instance, vocational robotics which are used to perform maintenance and construction tasks (i.e. repairs, inspections, cleaning, plowing) are driven by underlying economics which make them significantly more affordable than the alternative. This surplus value from the MUDs can be redeployed in a tax saving, new social service program or even simply in increasing the frequency of maintenance and construction, which in turn increases the functional accessibility of the sidewalk. Similar observations can be made for cargo delivery which can use efficiencies to reduce the needs for road maintenance and construction, offering opportunities for alternative programs or tax breaks, or simply to increase the affordability of home delivery which disproportionately benefits mobility reduced individuals. The key is ensuring that the potential detractors to accessibility that arise from the operation of MUDs are reduced as much as possible in at the design phase and deployment procedures, and that the surplus value created from their deployment is redeployed in the most constructive manner available.

Some thoughts for how to ensure that MUDs are a positive contributor to accessibility include sourcing labeled image data that specifically categorize edge cases of relevance to people with disabilities so the device can respond accordingly in a reliable manner. Should this be undertaken by a government, or consortium of interested stakeholders, this kind of labeled data could be offered up as a public good to other roboticists and road users. This would permit the benefits of accessibility-minded navigation systems to be widely scaled at no additional cost to users or manufacturers. MUDs can also integrate with other adaptive non-sensory technologies, such as Bluetooth Low Energy (BLE), to allow for non-sensory methods for MUDs to be perceived, with the MUDs effectively communicating their location to smart phones who can then transmit this to users through alternative feedback, such as haptic outputs. With a common BLE communication protocol for MUDs, it will be possible to ensure that the most comprehensive forms of accessible design are made available from the earliest outset, with the ability to be further built upon as systems improve. It is important that these kinds design approaches be sought out by manufacturers and governments alike, and that features of universal design be highly valued by both as well.

Bibliography

- Bakach, I., Campbell, A.M., and Ehmke, J.F. (2022). *Robot-Based Last-Mile Deliveries with Pedestrian Zones*. *Front. Transportation* 2:773240.
- Basecelli, Christopher. (2018). *Where the Sidewalk Ends and Robots Deliver: Setting a Framework for Regulating Personal Delivery Devices*. *Rutgers Computer & Technology Law Journal*.
- Brail, S., & Donald, B. (2021). *Robotic Cargo Transportation – New Technologies, Novel Practices, and Policy Readiness in Canada*. Transport Canada:
<https://tcdocs.ingeniumcanada.org/sites/default/files/2021-08/Robotic%20Cargo%20Transport%20-%20New%20Technologies%2C%20Novel%20Practices%20%26%20Policy%20Readiness%20in%20Canada.pdf>
- Brévignon-Dodin L. Regulatory enablers and regulatory challenges for the development of tissue-engineered products in the EU. *Biomed Mater Eng.* 2010;20(3):121-6. doi: 10.3233/BME-2010-0623. PMID: 20930319. <https://pubmed.ncbi.nlm.nih.gov/20930319/>
- Chatzimichali, A., Harrison, R., & Chrysostomou, D. (2021). *Towards Privacy-Sensitive Human-Robot Interaction: Privacy Terms and Human-Data Interactions in the Personal Robot Era*. *Journal of Behavioral Robotics*, 12: 160-174.
- Corno, Matteo., & Savaresi, Sergio. (2020). *Measuring Urban Sidewalk Practicability: a Sidewalk Robotic Feasibility Index*.
- Deng, Puyauan., Amirjamshidi, Glareh., & Roorda, Matthew. (2019). *A Vehicle Routing Problem with Movement Synchronization of Drones, Sidewalk Robots, or Foot-Walkers*.
- Dickinson, H., Smith, C., Carey, N., and Carey, G. (2018). *Robots and the delivery of care services: What is the role for government in stewarding disruptive innovation?* Melbourne: ANZSOG.
https://www.researchgate.net/publication/332857662_ROBOTS_AND_THE_DELIVERY_OF_CARE_SERVICES_What_is_the_role_for_government_in_stewarding_disruptive_innovation
- Hoffmann, T., & Prause, G. (2018). *On the Regulatory Framework for Last-Mile Delivery Robots*. *Machines*, 6: 33.

Jaller, M., Otero-Palencia, C., & Pahwa, A. (2019). *Automation, electrification, and shared mobility in urban freight: opportunities and challenges*. City Logistics.

Miller, Z., Kinder, J., & Moore, R. (2022) *Surface Robotics – Discussion Paper*. Institute on Governance.

Mintrom, M., Sumartojo, S., Kulic, D., Tian, L., Carreno-Medrano, P., & Allen, A. (2021). *Robots in public spaces: implications for policy design*. Policy Design & Practice:
<https://apo.org.au/sites/default/files/resource-files/2021-04/apo-nid311771.pdf>

Nelson (City) v. Marchi, 2021 SCC 41 (CanLII), <<https://canlii.ca/t/jjs98>>, retrieved on 2022-03-02