Envisioning Autonomy

by
Greg Martin

A thesis presented to OCAD University in partial fulfilment of the requirements for the degree of Master of Design in Digital Futures

Toronto, Ontario, Canada, 2022

Abstract

Autonomous driving is within reach. Machine vision and artificial intelligence are now mature enough to make driverless vehicle technology inevitable, but significant legal, ethical and design challenges remain.

Autonomous vehicles will offer precision, speed, and reduced reaction times, but these changes impact how pedestrians, cyclists, those using public transit and those with mobility issues navigate urban spaces.

Can the systems that will control networks of autonomous vehicles be designed to mitigate the safety imbalances that exist today for anyone trying to get across a city without a car?

What impacts will a network of vehicles that communicate invisibly and instantaneously have on the way we provide visual cues and signage to other road users?

When autonomous cars are inevitably tasked with choosing one life over another in serious accidents, are systems in place to evaluate bias in the algorithms used to make these decisions?

This project comprises discursive design objects intended to spur critical engagement with the legal and ethical ramifications of autonomous driving. By positing these issues in visually digestible form, the project aims to empower the public to better understand the implications of autonomous driving technology before its implementation.

Acknowledgments

This research was carried out and focuses on the ancestral and traditional territories of the Mississaugas of the Credit, the Haudenosaunee Confederacy of Six Nations, the Anishinaabe and the Huron-Wendat nations, who are the original caretakers of the land on which we stand.

I would like to express my deepest gratitude:

To my advisors Roderick Grant and Ali Qadeer, who understood the point I was trying to make even before I was able to articulate it. They challenged me to work harder, provided new ways of looking at and investigating the source material, and enriched my experience immeasurably.

To Adam Tindale and Christopher Moorehead, whose enthusiasm was infectious.

To Lloyd McArton, who gave me great research ideas even as I (accidentally) gave him food poisoning.

To my parents, grandparents and family for their support, and to my brothers, who are long overdue some credit for being the best siblings I could ask for.

And to my girlfriend, Serena, who supported me through a massive life change and kept our world running as I spent every waking hour at the laptop. I could not have done this without you.

Table of Contents

List of Figures

Introduction

We have been waiting for the autonomous revolution for decades. Since the earliest futurist tradition, the public has been captivated by the concept of people and goods using artificial, omniscient eyes and ears to move seamlessly between origin and destination.

Urban Canadians spend an average of 60 to 100 minutes on or near roadways each day; full-time workers spend nearly two hours a day on or near roads (Matz 216). The investment of so many hours into an often taxing task makes automation an enticing prospect.

Beyond time, safety is another concern; nineteen out of twenty crashes are caused primarily by human error (Singh 1). Perhaps the single greatest benefit offered by autonomous driving is the improved reaction times of machine vision systems, allowing for quicker classification of obstacles and improved collision avoidance.

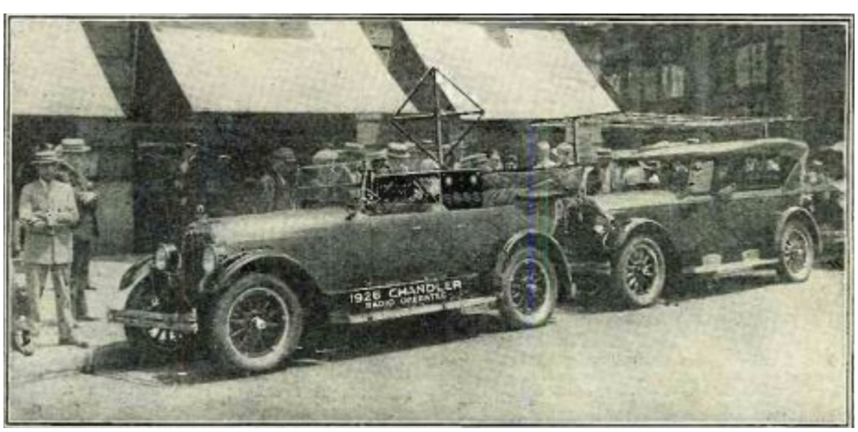

Figure 1 *American Wonder*, the first known driverless car, travelled unmanned down New York's Broadway in 1925, operated remotely via radio control (Goldhill).

Source: Wikimedia Commons https://upload.wikimedia.org/wikipedia/commons/d/d7/Auto1EMU.png

License: Creative Commons 4.0

Centrally managed fleets of autonomous vehicles also permit far more efficient travel in the aggregate, as techniques like carpooling and route planning offer an opportunity to put chronically empty passenger seats to better use. What autonomous vehicles offer is a usage model somewhere between private car use and public transport, a model in which vehicles group together like trips from otherwise separate vehicles and offer them to travellers at steeply discounted rates.

Even the most resolute skeptic would struggle to assert that autonomous driving offers no potential, but what progress has been made so far in the implementation of self-driving? It is an expensive proposition even to study, never mind bring to market. As of 2020, US$16 billion has been invested into the research and development of autonomous vehicles (Efrati), most of it shouldered by just a handful of deep-pocketed private enterprises.

In the public sphere, a lack of oversight over terminology and safety threatens to undermine the entire industry in its infancy. As governments and regulatory agencies scramble to keep pace with technological progress, iterative improvements in assistive driving technology are clothed in the futuristic verbiage of 'self-driving'. Features that automatically keep cars centred in lanes and assist the driver with parking are increasingly packaged and advertised as a type of autonomous driving, despite requiring constant driver supervision.

Elon Musk's Tesla is the best-known culprit of this kind of promotional conflation: its 'Full Self-Driving Beta' software package, while far from complete as an autonomous vehicle technology, permits its owners to use questionable and error-prone driver assistance features without any real oversight. In January 2022, more than 53,000 Tesla vehicles were recalled by the NHTSA in order to remove a firmware feature which, if enabled by the user, allowed its cars to roll through stop signs at up to nine kilometres per hour (Shepardson). While Tesla contends that this feature is only used when no pedestrians or other vehicles are in the vicinity, the shortcomings in its object classification software – which have resulted in the death of at least one driver (Boudette) – mean that

Tesla effectively cannot guarantee that their vehicles would stop this manoeuvre in the presence of pedestrians. The development, promotion and dissemination of software enabling the illegal operation of thousands of private cars shows how challenging it is for regulators to police over-the-air updates to vehicle functionality.

Multiple companies are investing billions not into developing an interoperable autonomous vehicle standard but into cornering and capturing the market for themselves. This arms race not only duplicates effort and cost but portends a worrying bias towards rapid development and release at the expense of safety and thorough testing. Decrying a lack of interoperability isn't just naivety either: the network advantages offered by a fully autonomous road system will require development of a robust open standard for communication and negotiation between vehicles, traffic infrastructure, and emergency services. Failing to design this from the start merely increases the wastefulness of the research process.

This project seeks to explore the following research question: *how can visual communication strategies replace informal patterns of communication and improve public safety in urban environments transitioning towards the use of autonomous vehicles?*

In responding to this question, I posit the following: already, the development of autonomous driving technology has become insular, privately funded, and neglectful of its own externalities. In developing clear visual standards for communicating intent and decision-making processes to pedestrians, driven vehicles and regulators, the entire industry will be made safer, fairer, and more efficient. By enforcing visual transparency, we can compel technical and ethical transparency.

Literature and Context Review

This project sits at the intersection of several modes of analysis, including urban theory and the politics of information. While it is not within the scope of this project to consider each fully, in reviewing the project's context it is important to give some sense of how autonomous driving threatens to upend existing relationships between citizens and the ways they get around.

It is necessary to define the terms *automation* and *autonomy* and help situate the specific concerns of this project within the wider implications of economic migration towards automation.

Establishing the influence that private vehicles have had on urban navigation design since their commercialization, we show how the influence of this infrastructure has inertia that extends into future design decisions, and how that inertia must be combatted to prevent worsening safety imbalances on urban roads. I demonstrate the implicit negotiation and mediation techniques that occur at today's intersections, and how those techniques are made explicit by autonomous vehicle infrastructure.

Finally, I make the case that full autonomy threatens urban navigators' basic right of access to information, and that our choices of how we collect and transmit data inform the accuracy and usefulness of the data themselves.

Automation and Autonomy

Automation is the reduction of human intervention, usually in systems of labour and production. Automation began with the mechanization of repetitive manufacturing and farming jobs more than a century ago. In more recent decades, automation has begun to replace more technical or precise jobs, especially in high-tech manufacturing, where minuscule fabrication tolerances are better met with robotic precision. As computational processing power increases and artificial intelligence gains the ability to make decisions based on complex and abstract inputs, more and more industries are exposed to this technological shift.

As such, one consequence of large-scale automation is structural disemployment, the obviation of entire roles and worker skillsets. Combined with the wide applications of a general artificial intelligence, very few jobs and skills remain safe from automation. Adam Greenfield's *Radical Technologies* identifies the opening salvos of this shift.

> Automation is a directional process whose initial stages we've already entered. In this respect David Graeber's empty, signifier-shuffling 'bullshit jobs' are a signal from the future. They're not so much a return as an anticipation of the repressed: the surfacing in the present, and pricing into contemporary ways of doing and being, of the recognition that there simply won't be enough meaningful work for anyone to do following the eclipse of human judgment (205).

Aaron Bastani calls the autonomous revolution humanity's third paradigm shift in *Fully Automated Luxury Communism*. The First Disruption, the domestication of animals and the cultivation of agricultural techniques, produced the conditions necessary for development: a general food surplus, the specialization of roles within a community and the collective living arrangements and commerce that specialization requires.

The Second Disruption, catalyzed by the invention of the steam engine in the mid-eighteenth century, again improved the yield of goods and food that a society could provide, laying the groundwork for improved life expectancy and universal literacy, as well as constructing the conditions necessary for long-term wage slavery, globalization and capitalism (9).

Like the first two, this budding Third Disruption is destined to further eradicate scarcity while, in current political conditions, exacerbating inequality (10-11). This project's research question offers a microcosm of Bastani's: will this revolution entrench current biases, or can the adaptations of the era also be used to remedy contemporary inequities?

A review of societal responses to earlier forms of automation suggests that capitalism does not provide the levers to meaningfully alter the relationship between labour and survival. Automation's potential to eliminate jobs was understood in the earliest political research on the subject – obviously, in many cases replacing comparatively expensive human labour is the point. Jeremy Rifkin's *The End of Work* draws a line connecting mechanization and the resulting disproportionate increases in Black unemployment to clandestine government attempts to quantify the negative social and economic impacts of automation. These attempts to use Black workers as unwitting subjects of economic study, made primarily in the mid-1960s, were quickly snuffed out by gathering enthusiasm for intervention in Vietnam: the engine of war reduced unemployment and boosted economic output. The commission's report, which called urgently for the creation of a universal basic income to combat an inexorable shift towards disemployment, was discredited as hysteria (81-83).

Rifkin also points out how the conditions of mid-century capitalism, especially the continued erosion of the power of organized labour, produced a scenario in which unions, not wanting to seem anti-progress, capitulated to the expansionary economic pressures of automation (84). Since organized labour has yet to recover even this level of influence, we should not count on its intervention in the years to come.

Autonomy and artificial intelligence are closely related, so it is important to first clarify definitions. While in a philosophical context we generally define autonomy to be an entity's ability to determine its own actions and fate, in technological discussions it usually refers to general, or strong, artificial intelligence. *General* artificial intelligence refers to an entity capable of judgment and analysis exceeding that of the human mind and remains as yet theoretical.

Narrow artificial intelligence, its more approachable counterpart, describes an agent that can complete a specific, defined and parametrized task. Narrow AIs require human assistance to define their parameters and provide training data and heuristics to evaluate the agent's success ("Strong AI"). Self-driving cars, and the wider class of autonomous vehicles (AVs), will use a highly sophisticated narrow artificial intelligence to complete the driving task at least as safely as humans in all conditions and environments. A condition in which an automated vehicle can achieve this is often referred to as *full autonomy*, but it does not constitute or require general artificial intelligence.

Hereafter we define *full autonomy* as the scenario in which all motorized vehicles employ Level 5 autonomy and human-operated vehicles are prohibited.

We also define *mixed autonomy* as the evolutionary phase in which some but not all vehicles in a city are self-driving, sharing the road with a plurality of human-operated vehicles.

The transportation industry is an excellent candidate for both automation and artificial intelligence. Millions of vehicles travel in lockstep and logjam all over the world along regulated, largely standardized roads. The driving task, though complex and demanding, is well-defined and precisely bounded. This differs from other tasks thought suitable for automation, such as shipping and logistics fulfilment, where fine motor skills and extremely accurate positioning and perception technologies are required. Some start-ups have made progress in this field, but it requires redesigning the picking workflow to accommodate the AI's limitations, a concession which is not an option for the replacement of outdoor, all-weather driving ("AutoGrasp"). This is a key problem facing the field: while some industries and workflows can be entirely redesigned to make better use of artificial intelligence, autonomous driving must outperform human drivers on their terms.

Onboard perception and processing of a vehicle's surroundings is feasible with current technology. It is what we might consider the low-hanging fruit of autonomy, by no means an easy task but perhaps the one with the best near-term return on investment. In fact, the transportation industry is already pervaded by components which use narrow artificial intelligence. As Ray Kurzweil put it (seventeen years ago) in *The Singularity is Near*, "if all the AI systems in the world suddenly stopped functioning…most transportation would be crippled" (289).

In terms of getting to full vehicular autonomy, it is easiest to start at the endgame and work backwards: imagine a world where the transportation of billions of people and tonnes of goods is achieved without someone ever touching a steering wheel or pedal. In this environment, all vehicles are constantly in communication with other vehicles, with the underlying infrastructure of a city, and with their passengers, providing efficient mobility for everyone, with the utmost safety and precision that sophisticated computer vision and machine learning can provide.

Of course, this scenario is unlikely to become a global standard. It isn't yet clear how widespread this type of navigation might be. Even within countries that choose to build out autonomous infrastructure, most experts don't believe that every area would be a suitable candidate for automation (Heinrichs 219); the dense urban corridors of cities and their suburbs are the focus of this project because they provide the necessary intensity to make automation economically worthwhile. What is clear is that most of the safety benefits of full autonomy only become possible when a critical mass of autonomous vehicles is on the road. Over time, even the most alert, cautious human driver will be unable to match an autonomous vehicle for speed of reaction time and mental stamina, and we should expect society to enthusiastically support the transition to full autonomy as the technology matures.

So what happens until then?

The proposition of governments mandating the use of fully autonomous vehicles would, for most, constitute civic overreach; it would only be palatable once the safety benefits of fully networked automation become indisputable, and would likely be phased in over a number of years. In the meantime, how will we navigate the liminal space between the cities of today, choked with human-driven vehicles, and the utopian promise of full autonomy?

Like with all exponential growth systems, the inflection point of the technology – what Ray Kurzweil calls the 'knee' of the curve – will seem distant until the moment of its arrival. Although retrofitting or replacing equipment for autonomous operation is a significant expense, human capital will quickly become more expensive than its automated replacement. Paul Romer points out that the price model collapses for these sorts of technological advancements: once the fixed cost of the change is recovered, marginal units are available at effectively zero cost (3). Once fully autonomous trucking becomes feasible, regulated, and safer than manual driving, the industry will transition quickly to a more efficient labour pool of tireless self-driving vehicles.

These changes will be swift, and they will catch many by surprise. Although the technology is still in its infancy, just as with the rapid expansion of Internet technologies and its secondary industries there is a real risk that its progress outpaces our ability to regulate and thoughtfully respond to those who stand to make their fortunes from its hastened introduction. An informed public is one able to advocate for its own safety and protection. This project begins the conversation on how we might best reach out to people to inform them not only of the positive potential of autonomous driving and automation, but of the ethical and economic concessions they may be asked to make in the coming years.

Cars Come First

Is a city of automated vehicles likely to look or feel significantly different?

Cities adapt to new methods of getting around all the time. Many airports have now cordoned off waiting areas for the specific logistics of rideshare app pickups and drop-offs. Some North American cities have made investments towards cycling infrastructure that emulate more significant shifts taking place in continental Europe.

The scale of these upgrades, however, pales in comparison to the task of preparing our existing road networks for autonomous use, road networks that are already in need of revitalization just to maintain current usage patterns. Proponents of these changes are likely to say that automated cars and buses will become just one more type of vehicle, sharing the road with other transportation. The more likely truth, however, is that migration towards autonomy will exacerbate the current privileging of privately-owned cars.

The advent of safe, autonomous movement by car is likely to increase overall travel demand, at least in the short term. Non-drivers, children, and the elderly will be provided a method of direct transportation previously inaccessible to them. For existing drivers, the onset of autonomy will not in itself reduce travel demand. The ease and novelty of the technology will likely induce additional travel by all able who are able to access it in its early years.

The resulting change we might expect of most cities, then, is to provide for more cars by adding more roads. We already know that the strategy of expanding road capacity to alleviate traffic is a myth; new lanes almost universally induce additional demand and additional congestion. As Jane Jacobs pointed out sixty years ago, "if vehicular traffic in cities represented some fixed quantity of need, then the action of providing for it would produce a satisfying and fulfilling reaction" (353). Despite thorough research into the phenomenon of induced demand, road expansion continues in many cities.

The increase in demand precipitated by full autonomy, however, does not need to be met with more vehicles or more roads. An alternative way to satisfy increased driving demand is the implementation of shared ownership models. A fleet of autonomous vehicles can be managed and maintained by a central entity, with vehicles themselves operating round the clock. Autonomous vehicles are uniquely equipped for the success of such a program, and the increased efficiency of shared fleets could allow the travel needs of a community to be met with fewer vehicles; today, the average car is parked for approximately 95% of its lifetime (Barter). A Hungarian study explored the feasibility of reclaiming some of Budapest's parking spaces under a variety of vehicle ownership scenarios. Key to the benefits of shared autonomous fleets is the vehicles' interchangeability: summoning a specific private vehicle requires a clear path from every car to the lot's exit, but standardized fleet cars can park themselves into much tighter grids and simply assign the closest available vehicle to a trip request.

There are additional benefits to the shared ownership model. The Shared Mobility Principles for Livable Cities – an initiative started by Zipcar.com co-founder Robin Chase and with signatories including Uber, the City of Paris, and the C40 city collective ("Signatories") – lays out the argument for a mandatory shared ownership model as being safer, more accurately priced, and more equitable:

> A world in which every individual in dense urban areas owns their own AV would distort land use and other behaviors and dynamics in ways that would be severely detrimental to cities...When a vehicle is shared, the decision to make a trip includes the full costs (depreciation, insurance, maintenance, as well as refueling) and therefore sets a higher cost hurdle rate to make a trip. In dense metro

> areas with scarce road resources, we want to discourage low-value car trips and therefore more efficient use of space…Initially, AVs will be very expensive. When AVs are in a shared fleet, the benefits can be purchased at low cost – by the seat for a specific trip – rather than having to purchase an entire vehicle ("Why Principle 10").

Regardless of the ownership model, the study found that autonomous vehicles show significant potential for the saving of space. Even in a private ownership scenario, AVs' ability to drop passengers at their destination and park themselves a significant distance away to await summoning means that, certainly in cities, any kind of autonomy improves the use of space (179). Dirk Heinrichs estimates that up to 60% more parking spaces can be provided in the same square footage thanks to the precision and density of autonomous parking (222).

Despite their ability to save space and the improved fuel efficiency that is implied by their arrival to market amid a global transition to electric vehicles, autonomous vehicles are likely to increase emissions: given the increased demand for vehicles, the embodied carbon of car tires, and the external carbon emitted as part of much electricity generation, continued car dependence, even with autonomous vehicles, is likely to continue putting upward pressure on carbon emissions (Eyre 48). For all of the vehicles' positive externalities, cycling, walking, and public transit remain the modes of transport that must be incentivized in the next century.

Personal vehicles, however, will continue to provide a significant portion of overall travel for the foreseeable future, especially in North America. From that perspective, we can be grateful for the introduction of autonomous vehicles, which improve on contemporary cars in so many ways that it may feel disingenuous to compare them. They will be safer, more efficient, likely electric, and they will free up countless hours currently given over to the driving task for business or leisure, significantly increasing overall productivity. They will, however, still use roads.

Today's roads require additional space to allow for driving idiosyncrasies and human reaction times. If autonomous vehicles were invented and distributed in a world that had never seen manually operated vehicles, we would design our roads differently. They would likely be narrower, and with fewer lanes. Autonomous cars drive closer together, more quickly, and are in constant communication with one another. The amount of land given over to roads and parking facilities could be drastically reduced.

In the real world, however, the car is king, and has been for over a century. All other forms of urban transportation defer to the mobility, ease, and ubiquity of private car travel. This prioritization, far from being some natural equilibrium, is the result of stakeholders' concerted interest in the continuing dominance of cars, namely the automotive manufacturers and fossil fuel companies that powered America's post-war economic engine.

To see thumbs on the scales of infrastructure investment, one need only review the powerful advisory positions given to automotive executives in the leadup to Eisenhower's spearheading of a trans-national federal highway system. This enormous public works project, paid for largely by consumers through a tax on gas and steered through Congress by a team including General Motors CEO Charles Erwin Wilson (Stromberg), was the beginning of a North American tendency towards heavy road usage, urban sprawl, and lower density.

(The interstate system being largely funded by gasoline taxes recalls the self-serving machinations of the very first Michelin Guides, formulated to induce people to treat restaurants in nearby cities as destinations, intensifying tire wear and increasing sales.)

This inclination towards expanded private car travel was marketed using immersive, futuristic design objects not unlike the discursive elements of this project. These slick, expansive exhibitions provided a highly curated, don't-look-too-close glimpse at the future offered by cars and the roads that carried them.

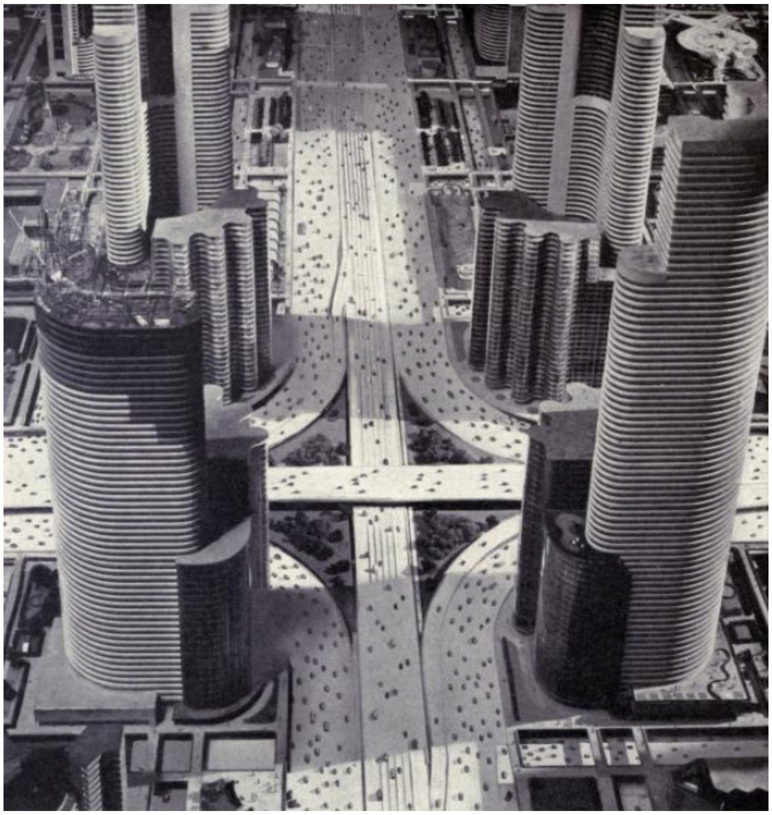

Figure 2 A portion of *Futurama,* a one-acre diorama designed by Norman Bel Geddes for the 1939 New York World's Fair. The exhibit featured over 50,000 moving vehicles.

Source: Wikimedia, https://commons.wikimedia.org/w/index.php?curid=67871017
License: Creative Commons 4.0

Futurama was a one-acre diorama designed by Norman Bel Geddes for the 1939 New York World's Fair. Sponsored entirely by General Motors and exhibited in their 'Highways and Horizons' pavilion, the enormous exhibit is credited with introducing American audiences to a nationwide 'expressway' system for the first time. Visitors enjoyed the installation by use of a slow-moving, comfortable and compartmentalized monorail which floated above the diorama below. As Douglas Adams of the Rhode Island School of Design wrote of the exhibit:

> The effect of the Futurama was created by a number of environmental factors: sound, light and colour, all interacting to simulate an airplane flight. This 'synthetic space' was essential to the visitor's experience of a 'magical flight through space and time'. The term synthetic here implies a space designed to influence behaviour...The teardrop vehicles exhibited on the proposed roadway, however radical

> in form, reflected the assumption that the private automobile would continue as the dominant means of travel (22).

The choice to have visitors to the exhibit enjoy the diorama from above is certainly a purposeful one. As Stephen Graham notes in *Vertical*, "the use of flyovers, air travel, skyscrapers, cable cars, elevators and Google Earth...produce powerful new aesthetic sensibilities of megacity life...elites now consume the city as an aestheticized landscape from above. The notorious 130-mile traffic jams in rush hour...the snail-like progress of the city's totally inadequate public transport system...all are rendered as little but a backdrop" (102).

American automakers had a guiding hand on the rudders of post-war expansionary policy, paving the way for automotive domination. The effect of 'gliding over all' is an intoxicating one, a way for viewers to distance themselves from the monotony and bureaucracy of reality. The diorama and its modern analogues are effective ways to instil within the average viewer an intense sense of future. The unreachability of utopia and the fanciful sensation of flying over such a space went hand in hand. Architectural historian Adnan Morshed writes:

> The simulated voyage over the Futurama was intended to resolve the optical limitation of earthbound views and, more important, the philosophical problem of experiencing the utopia that, as an ideal condition, eludes us in reality...Seeing the utopia – which perpetually eludes our epistemological boundaries – required, here as in the Futurama, an equally fictional mode of spectatorship (78).

Despite the success of that World's Fair, particularly the General Motors pavilion, few would argue that Bel Geddes' promised utopia arrived, even with the delivery of thousands of miles of paved expressway. As architecture critic Paul Goldberger writes, "Bel Geddes figured out what the modern American city was going to look like before anyone else did. He knew that visual excitement and energy had to be a part of modernism. He just failed to understand that there were other things that made cities work, such as streets and neighbourhoods, not to mention serendipity, and that these things were incompatible with the world of the automobile" ("Back to the Future").

Today the bar for futuristic experiences is set high by modern entertainment and computer imagery, but we see glimpses of this attempt to influence consumption and legislation by eliding detail, complexity, and nuance in promotional videos for Elon Musk's Boring Company in which we watch, from the same top-down perspective, gleaming Teslas slip seamlessly from congested roadways into a futuristic network of subterranean tunnels (via means of an unmarked, uncovered, car-sized hole in the ground, into a tunnel whose only real-world exemplar has no emergency exits) ("Boring").

The inertia of this infrastructure should not be discounted: once the system is built, it becomes nearly impossible to conceive of another way of doing things. Such was the scale of initial investment required by the government that it would be political suicide to renege on this vision of a highway to every city, even as the infrastructure itself crumbles. The American Society of Civil Engineers estimates that, as of 2021, over 40% of the American road system is now in poor or mediocre condition ("Roads").

This inertia is what we are up against: the costs of doing nothing are invisible, a slow leak behind walls of maintenance costs and consulting fees, while the upfront costs of rebuilding or overhauling these infrastructure systems seem staggering by comparison. Nevertheless, it is this work that produces roads and urban spaces that address the needs of all.

Discursively this project functions as a kind of diorama, one that eschews a utopian overview for a realist cross-section, seeking not to ignore the complexity of life at ground level but splay it out for analysis.

Negotiation Tactics

While the number of passenger deaths in vehicles has dropped 25 percent since 1975, the number of pedestrians killed in the US hit a thirty-year high in 2019 (Beresford). In the five-year period ending in 2020, more than 1060 pedestrians or cyclists were seriously injured or killed on City of Toronto roads ("Vision Zero"). More than three quarters of those pedestrian fatalities are the result of high-risk mid-block crossings and left turn and right collisions at signalized intersections ("Vision Zero 2.0"), areas of informal engagement between pedestrians and drivers. These accidents are the result of policy, of an urban navigation network that accepts a variety of travel types but fails to provide equitable access to their safe use.

Our roads are communal places, forums where vehicles of various sizes meet and negotiate to find a mutually agreeable set of outcomes: that every vehicle arrive at its destination in as timely a fashion as possible. This process of negotiation occurs at every intersection, on-ramp, roundabout, and driveway.

Most people can conjure images of large, unmarked intersections overflowing with vehicles, each car seemingly facing a different direction, slowly honking and nudging along the tarmac to its destination. In Toronto, the decisions of each entity within the navigation process are largely mediated by types of infrastructure: the physical infrastructure of traffic signals and signage, the social infrastructure of etiquette and caution, and the legal infrastructure of regulation, insurance, and enforcement.

With each trip we make, we are processing a dynamic urban environment and completing hundreds of tiny negotiations with thousands of unwitting parties. If these negotiations are so marginal and so common, how do they almost always proceed as planned, and how do we all get where we're going?

To understand this implicit conversational layer, we must make use of concepts from the field of wayfinding. Wayfinding is generally defined as the set of methods that individuals use to navigate a city using its paths, edges, and landmarks. This definition, described in Kevin Lynch's *The Image of the City*, was expanded upon in later years by environmental psychologist Romedi Passini to describe the ongoing dialogue between navigators and the stimuli provided by their immediate environment (29).

This dialogue is one that any navigator has during a trip through a city, not just with the built environment around them but with all other nearby vehicles and individuals. Each 'line' of dialogue, each discrete scenario with its own set of outcomes, is a *decision point*. A decision point may be a fork in the road, an office building's directory, or a ramp next to a flight of stairs. Navigators are constantly arriving at and evaluating decision points and attempting to select the best outcome based on the available information. It may not be discernible which outcome is best in a situation, so the task of wayfinding design is to make it clear to the navigator what their options are and, if possible, which selection is optimal.

Resolutions to decision-making processes that are not necessarily optimal but meet some minimum criteria are referred to as *satisficing*, a portmanteau borrowed from behavioural psychology. A satisficing decision is perhaps not the optimal one, but both *satisfies* the navigator and *suffices* for achieving the task. For an individual trip this may mean finding the fastest possible route to a destination within five minutes of route planning. A slightly quicker trip may be possible, but the time and effort investment required to find this route is excessive. Each person's criteria for a satisficing plan may differ; some want to get there quickly, and some want to get there cheaply. There exists for each person an intuitive frontier between cost and speed that confirms a route plan as being 'good enough', and this frontier shifts in the individual according to their circumstances of travel. Planning a road trip weeks in advance, for example, has different satisficing calculus than a frantic GPS search for the nearest hospital.

Satisficing outcomes are crucial to the operation of complex transit interfaces. Not everybody gets where they would like to go as quickly as they would like to get there, but there is a general sense of viable compromise. Like any negotiation, each party has a degree of leverage that they may use to influence the decision-making process. This leverage can be perceived by other parties as less or more influential than it may actually be, and usually a party's perceived leverage impacts its priority more than its true underlying power. We think of this as bluffing in certain environments, but in our urban environments it is more of an evolutionary risk avoidance tactic. Even a cautious, defensive driver in a three-ton Hummer is likely to receive deference from pedestrians more often than an aggressive driver in a hatchback.

At a signalled four-way intersection, these negotiations are mediated by traffic lights and crossing signals. Each party arriving at the intersection enters a priority list that is resolved by mutually agreed-upon rules which are codified and symbolized with signals and lights. Perceived leverage exists in these situations too. Cars, trucks, and buses can nudge into crosswalks and speed through 'stale yellows' by virtue of their size and momentum advantage over other parties at the intersection. Each skirting of intersection etiquette is a 'who's going to stop me?' to comparatively weaker co-negotiators. As mentioned above, similar creeping infractions at right turns are responsible for a quarter of pedestrian fatalities in Toronto in the last five years.

Consider another example: a wall-to-wall traffic jam on a busy highway. Three lanes of traffic are filled and virtually stationary, but the emergency shoulder to the left remains open. Nothing physical exists to prevent vehicles from using it, but the solid white line on the tarmac acts as mediator, signalling to all parties its ineligibility for use. Mediators only work, though, while they hold the confidence of all parties involved. When a vehicle speeds down the shoulder to the anger of drivers nearby, the social contract breaks down and passive mediation becomes ineffective. Other drivers may follow, and soon a single act has changed the rules of the road.

In some respects, the prospect of full autonomy offers a solution to these situations. A road system filled with vehicles connected to a central infrastructure can be compelled to follow traffic signals and provide safe passage for smaller, more vulnerable modes of travel. A system like this does, however, require a completely different class of negotiation to those occurring at today's intersections. Nobody gets out of their cars at a crossroads and haggles as to who should go first, but some of the approaches to autonomous intersection management are precisely that: protocols for a kind of instantaneous haggling.

Stefano Mariani produced a taxonomic survey of proposed approaches for intersection management and traffic flow optimisation in a full autonomy scenario. Intersection management is identified as a resource allocation problem, one in which "vehicles compete for the right to acquire exclusive access to the portions of the intersection they have to cross" (10).

Within this problem space there are many available approaches, each affording varying degrees of freedom to vehicles in the intersection. A *centralized* approach requires each vehicle approaching an intersection to register with an external coordinating agent, who provides motion instructions akin to air traffic control. *Negotiation*-style protocols enable vehicles at intersections to communicate and, in effect, jockey for position in the priority list depending on vehicle-level factors. *Emergent* protocols for intersection management take inspiration from nature and game theory, proposing that vehicles do not explicitly coordinate but rather act in a purely reactive way, using either a kind of implicit perception of the intersection's state or a system of "virtual pheromones" (10) to determine the intent of each vehicle.

There are downsides to each approach. For centralized protocols, the physical overhead and basic computational requirements of central intersection management may not scale to an entire city. For negotiation protocols, the prospect of introducing a mediation currency, either real or virtual, to resolve priority raises ethical issues. In high-traffic scenarios, vehicles may be priced out of crossing an intersection entirely – what Mariani refers to as "starvation". This implementation suggests the possibility of having to abandon your rental at a busy intersection because you cannot afford the 'micro-toll' to cross in front of wealthier traffic. For emergent protocols, the total lack of explicit coordination between vehicles hinders the system's ability to guarantee safety or timely crossing for all parties. It is most likely, Mariani concludes, that intersections will dynamically adjust their management according to traffic levels and extenuating factors such as the presence of emergency vehicles (25).

Regardless of their individual quirks, all scenarios suffer the same horizontal drawback: these approaches only function insofar as every party that seeks to use an intersection is connected to the allocating infrastructure, i.e., is a fully autonomous, self-driving car. In a mixed autonomy environment, at least some of the vehicles jockeying for position at an intersection will be human-operated. This may mean that these resource allocation systems cannot be instantiated until the compulsory introduction of self-driving vehicles.

Even then, not all parties at an intersection will be connected to its infrastructure. Pedestrians and cyclists are still not accounted for. The utopian goal of intersections without signals, of latticed streams of cars interweaving at dizzying speeds, seems unattainable. How can we include and protect human navigators in these scenarios?

A clear system of visual communication remains the primary strategy for integration of human navigators into mixed and full autonomy transit environments. Both within and without these vehicles, the safety of an intersection depends heavily on the mutual trust that negotiators place in each other, a trust that is derived almost entirely from visual cues: a vehicle in motion towards you, a quick double-flash of headlights letting you go ahead, a sweep of the hand from a pedestrian crossing ahead of you. As this vernacular vanishes, it becomes vital to codify the ways we signal intent and momentum at intersections.

Who Wants to Know?

Our world is permeated by dozens of data layers. Some of them we use in moving around our community, searching for information, or communicating with others. Some are less useful but no less ubiquitous: Pokémon Go gyms (Greenfield 141), Snapchat maps of cultural hotspots, and city sensors reporting an inch of snow or a tipped-over garbage can. Others still are totally hidden from us, encrypted, garbled or unintelligible.

We do not speak binary. Some interlocutor must exist between the trillions of zeroes and ones we generate and the information they represent so that we can derive meaning or utility from their content. Any data we read as an image, an icon, or a line of text is a projection of the real thing, a shadow of a dataset cast onto the backdrop of the physical world. Like real shadows, in viewing these projections we grasp the outlines but not the depth or detail of the objects they describe. As consumers, we have little control over those projections, but as more and more physical processes are sublimated into data exchanges, the right to define the shape and extent of those projections becomes paramount.

Other data layers exist in a more abstract sense than zeroes and ones. On the average neighbourhood high street, each store's opening hours, their address, and their wares are broadcast physically to passersby using signage or other visual cues (a door left ajar, an unlit storefront), but they also exist as intangible sets of information known to that neighbourhood's locals and visitors. If a store neglects to post its opening hours in the window, those opening hours do not cease to exist. For trendier bars and restaurants, the omission of information from the public square can be part of their appeal, a deliberate shrouding from view that invites exploration and conveys cachet.

Everything in a shop window is a projection of that shop's data into the physical world, and we have settled on acceptable standards for providing that information. Other data still might be widely known about a store – its owner's reputation, its usual temperature, whether its wages are fair – without being provided as part of the public-facing presentation. The display of some data, such as occupancy limits or food safety inspection results, are even mandated by municipalities. The question, then, is one of deciding which projections matter and, of those, which must be shown compulsorily.

Data layers are not flat planes extending out in all directions. Data has density and scarcity, thickness and flimsiness. It is a viscous fluid, one that collects in busy areas and dries up in a city's fringes and capillaries. Interference is frequent in urban cores: most urban navigators know instinctively not to trust their GPS location when surrounded by skyscrapers or turning down a side road ("GPS Accuracy").

And since data of this kind is not free to collect, it tends to pool where its value is greatest. Data's depth and granularity is dictated by capitalistic mechanisms. In one Louisiana town, Google Street View mapped only the wealthier areas and ignored the poorer neighbourhoods (Sutter), and while it undoubtedly furnishes an essential resource, the mapping system has a problematic relationship with consent. While it is relatively simple for someone to request that Street View blur out their house or car, it remains at heart an opt-out process, one that privileges those with the time and resources to submit a claim. Google counters this with the argument that its service only shows public streetscapes and is therefore within the law (Kopytoff), but it stretches credulity to think that by engaging in visible activity in a public space a person consents to its being pictured, uploaded, and forever preserved in an offshore data centre.

The misconception of data as an inherently precise and dispassionate medium persists today, driving civic decision-making and resource allocation. Cities tout the introduction of networks of Internet-connected sensors as watershed advances in urban monitoring and data analysis, but as Adam Greenfield points out, "however thoroughly sensors might be deployed in a city, they'll only ever capture the qualities about the world that are amenable to capture" (52). There exists a kind of white coat effect when it comes to data, a sense that it must be an objective representation of the thing it purports to summarize: "The authorship of an algorithm intended to guide the distribution of civic resources is itself an inherently political act...nowhere in the extant smart-city literature is there any suggestion that either algorithms or their designers would be subject to the ordinary processes of democratic accountability" (59).

The data layers we use and produce are essential utilities that are yet to be supported, regulated and protected as such. Citizens are alternately active and unwitting contributors to these datasets, and recent history is littered with examples of that data being monetized, poorly secured, or compromised without our consent.

Like all utilities, access must be guaranteed regardless of relative expense, and strict standards must be put in place to protect and ensure its accuracy. Such shortcomings are also warnings of the widening gap between the data-gathering proficiencies of large technology companies and the oversight capabilities of governments responsible for ensuring fair play.

The politicization of efforts to enshrine into law principles of net neutrality – the set of rules ensuring that wired and mobile broadband providers treat all data transmissions that go through their pipes or airwaves equally, irrespective of content (Werbach) – demonstrate the uphill struggle that faces any attempt to guarantee access to these emerging resources. In 2017, the Trump administration, led by FCC chairman Ajit Pai, began the process to roll back protective policies enacted during the Obama administration, citing a need to "'modernize [policies] to match the reality of the modern marketplace'" (Brodkin). One of the Biden administration's first executive orders was to reinstate the rules rolled back by the Pai-Trump FCC, but future Republican administrations seem likely to continue a regulatory tug-of-war that precludes necessary stable long-term investment in public broadband infrastructure.

While the public safety implications of autonomous driving may introduce bipartisan urgency where net neutrality has fallen short, the extent of corporate regulatory capture and the growing technical illiteracy of many governments suggest that an arduous process awaits.

Our roads and traffic networks require and produce enormous data layers, which are being continuously tapped into and contributed to by drivers and their vehicles. A speed limit exists for every road in the country. For motorists turning onto an unsigned road, this limit is provided contextually, either by the average speed of the traffic they're merging with or by subtler cues such as the number or width of the road's lanes and medians. Increasingly often, in-car GPS systems and smartphone applications furnish drivers with that information, reducing the demand for physical signage.

Stephen Graham's *Vertical* contextualizes these data exchanges as negotiations between different vertical layers of the atmosphere, from the military-controlled upper ionosphere all the way down to the city infrastructures of subways and sewers. The top down, bird's-eye nature of so many of today's visualizations of cities compresses the vertical plane flat, eliding conversations about equity and control that can occur within the same apartment building. Even the increasingly crucial statistic of elevation above sea level is lost in this perspective.

This project suggests we transpose that metaphor. City thoroughfares comprise horizontal segments of space, like slices in a loaf of bread: pedestrian space on sidewalks, vehicle space in roadways, transit space in high-occupancy lanes. Within those slices, the quality and communicative efficiency of data varies wildly. A motorist in a relatively modern car has access to orders of magnitude more data concerning their environment than does the pedestrian alongside them. Motorists are fed a vast amount of information on large, rich displays and interfaces. This is a conciliatory act: such is the potential for destruction in these vehicles that drivers must be physically and informationally centred as the primary users of any space. (Even the busiest pedestrian crossing occurs at a red light. At a green light, the intersection resumes its normative state: prioritizing the safe passage of cars.) These imbalances of access to data are the ones that autonomous vehicles have the most potential to reinforce.

As these features become commonplace, we might expect that cities and municipalities will commit fewer resources to signposting speed limits. Signs are expensive, imperfect, and not easily changed; if a speed limit changes on an unsigned road, updates are as simple as uploading new data to navigation databases. Physical signs need to be replaced or amended, and this change itself often needs its own sign warning of a new sign ahead.

The evolution of modern driving has tended towards fewer physical interventions, uploading responsibility to increasingly sophisticated in-car systems. For drivers with older or simpler cars lacking these modern conveniences, their ability to tap into these abstract data layers is increasingly impacted by the prevalence of more advanced vehicles. It is possible to envision a future scenario where the runaway pace of dynamic in-car 'infotainment' technology halts maintenance on expensive, static physical interventions.

The question becomes one of equity of access to information. For the interregnum between widespread autonomous vehicle adoption and federally mandated full autonomy, a dwindling number of manual drivers will attempt to navigate cities that are increasingly hostile to their presence. We should not conflate manual drivers and those without access to current-generation GPS technology; any modern smartphone can achieve roughly the same functionality as an integrated navigation system, so at present, older car owners can easily access the same information as new car owners, but a modern smartphone should not be a minimum barrier to entry for drivers. Once the technology provided by a newer car is not just augmentative data, but autonomous software core to its safe operation, we risk pricing out a significant portion of motorists from roads.

Perhaps we should rely instead on context and trust the judgment of our fellow motorists. The removal of white lines from the middle of roads is associated with fewer accidents and slower driving speeds, as the uncertainty introduced by 'naked tarmac' produces cautious decision-making (Beresford). The problem is, again, forsaking a minority group in the interest of most drivers. If cars were the only vehicles that ever touched roads, perhaps this would be feasible, but the modes of travel that we truly need to incentivise in this century – cycling, walking, and public transit – require additional visual cues to stay safe. Frictionless simplicity is not an option.

In *A Short History of the Printed Word*, Chappell and Bringhurst discuss the productive function of friction in technology, and about what its progressive disappearance means for the process of production. This helpful friction is manifest both physically and culturally. A very minor change in nib shape from the broadpen to the pointed quill in the 16th century brought about significant changes in handwriting. A century earlier, the Church was threatened by the democratization of the printed word that Gutenberg's press provided and sought to discredit such a medium as artless and glib. Aaron Bastani writes that "while the Gutenberg press was

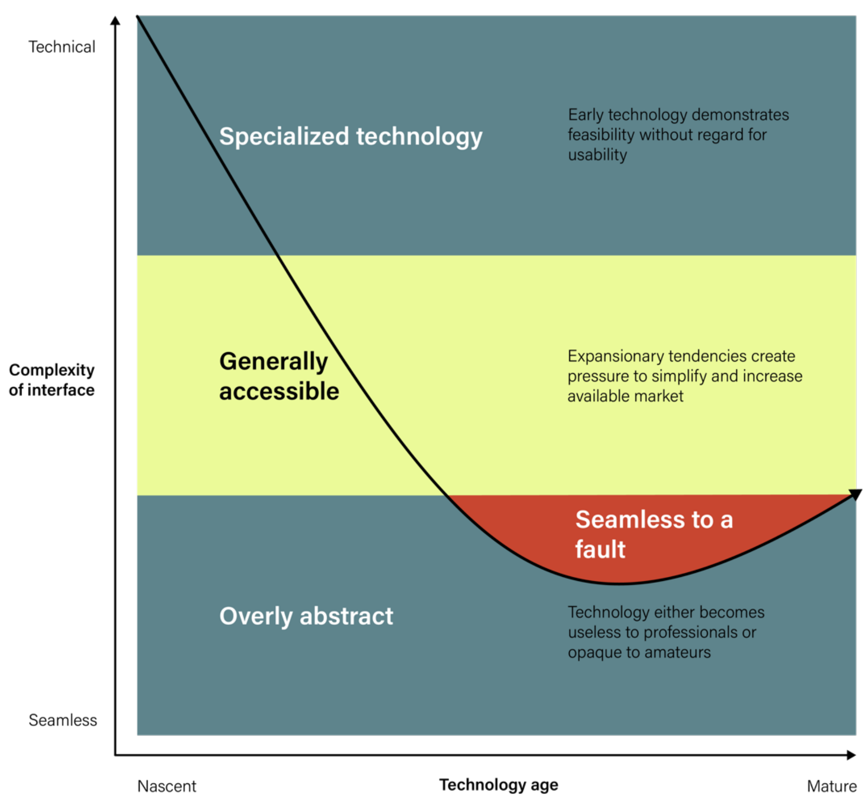

Figure 3 A diagram demonstrating the impact of market growth on a technology's interface.

profoundly disruptive, it only led to social transformation once it became so mundane that a little-known theologian could have his ideas printed by people he had never met and, in a matter of months, discover an audience of millions" (241). This same technology was eventually retired in favour of more frictionless digital typesetting, a change which resulted in the disappearance of entire classes of jobs.

This is the life cycle of technological friction (see Figure 3). Frictionlessness has a democratizing effect in the early stages of a technology's development, lowering barriers to access and making its use easier. At a certain point, however, the search to make something easier to use involves the incorporation of some technological element that obscures the tool's operating mechanism behind some layer of abstraction.

In software design, this can mean a simplified, 'streamlined' interface that alienates newcomers and frustrates professionals. For delivery apps, this layer of abstraction is the app itself, obscuring the entirety of the labour and goods exchange involved and granting users a blissful ignorance as to the externalities of their order. In the case of autonomous vehicles, their breezy, seamless, seemingly free travel threatens to crowd out pedestrians, cyclists, and other road users and reduce our obligation to provide them a safe medium.

Once this abstraction occurs, the reduced friction causes reduced access, access that privileges high-information or high-status users. There exists an equilibrium for each technology between ease of use and ease of access, and for public technologies such as transportation signalling, this equilibrium must be fixed and adhered to.

Culturally, workers and creators in any industry who hearken back to pre-digital processes as a production method are often heralded as 'true craftspeople', seeking immediacy and tactility. This is a manifestation of the process described above, of the pendulum swinging past its equilibrium before reverting to the mean. The resurgence of celluloid film technology in independent filmmaking is testament to this process: at a certain point, digital filmmaking loses something essential to the process. Even digital shots made to look like indistinguishable from film are denigrated as simulacra, signifiers of a director unwilling to 'do things properly'.

If this line of thought sounds familiar, that might be because it's everywhere. Even the plowhands of *The Grapes of Wrath* attribute the growing malaise among small holding farmers to the mechanization of the harvest, and to the attendant accumulation of resources and rentseeking that such mechanization enabled. Many characters see the Oklahoma dustbowl that ravages their land as spiritual retribution for failing to sufficiently engage with the soil that provides their livelihoods, and consider the less gruelling option of mechanical harvest as false (35).

Friction, then, does not mean an obstructive or counterproductive force in this context. It is more of a lack of self-effacement, the act of broadcasting information about a process, about the work that goes into its creation, about how to use something, that remains so vital and is at such risk of vanishing in a data-driven world. With the act of driving nearing its technological endpoint, we must be wary of crossing into an area of frictionlessness that obscures rather than elucidates.

In terms of establishing a fair baseline for visual communication during an autonomous revolution, we should take many of our lessons from the fields of visual communication developed over the last century. Environmental graphic design provides an extensive precedent for how people perceive, process and act on information. It is not enough to know that autonomous vehicles are in seamless, instantaneous communication with one another. We need to ensure that they transparently communicate with everybody else.

Methods & Methodologies

Discursive Design as a Methodological Framework

This project is constructed using principles of discursive design, a term popularized by researchers Stephanie and Bruce Tharp ("Mode" 406).

Discursive design is a field of design contending that design can function entirely in its ability to produce a discourse, a discussion of aims, goals and challenges on an individual, community, or societal level. "Discourse is why [a discursive object] exists." ("Discursive" 77). It is well-suited to complex problems such as these because discursive design "uses its tools to affect reflection, acknowledging and trying to unpack the complexity as a means of possibly progressing toward a preferred state, or at least identifying attributes of what one might look like or not look like" (78).

Discursive design is heavily categorized and subcategorized, at least in the Tharps' formulation. They position discursive design as the fourth of four fields of design practice, the others being commercial (design for profit), responsible (design for equity), and experimental (design for process). Successful discursive objects "essentially hinge upon the designer getting the audience to reflect upon their discourse" (112).

Discursive design is further categorized across two axes: internal-external, where the audience for the outcome is either within or outside the field of design itself, and terminal-instrumental, in which the object of the design process either is the intended goal of the design process (terminal) or is instead a part of other design activity whose goal is larger than or separate to the design process of the thing itself (instrumental).

In this discursive design taxonomy, the project outcomes are *instrumental-external*: the intended audience is outside the field of design, and the design object is a part of a wider, as yet incomplete, design process. The project is intended for a general audience, not as a means of publicity but as an imperative: it considers the existing failings of urban public communication and how autonomous driving may exacerbate them. Improving public awareness of the governance and ethical issues surrounding the onset of autonomous driving is an externally facing project.

By the same token, the specific outcomes demonstrated as part of the project communicate possible directions that regulatory bodies could take to protect public access to safety information, and not those regulatory notices themselves. It is not within the scope of this project to propose environmental graphic design standards for physical media in an autonomous driving environment, but it is within project scope to point to their absence and the implications of that absence.

Discursive design makes heavy use of the concepts of *friction, dissonance* and *scenario* as factors in the dissemination of ideas.

Friction and dissonance are closely related concepts in this context. Much discursive design co-opts aspects of what it depicts as a visual shorthand to get the audience on the same page. For example, a basic white page of unlabelled ideograms demonstrating how to assemble a teleportation device with only a hex key uses the audience's collective understanding of an existing good – in this case, an IKEA manual – to get them 'up to speed' in as little time as possible and ready to engage with the discursive matter at hand, the teleportation machine. The discursive designer "games the normative system" (13) by using our existing relationships to objects that are similar to the design artifact to convey complex ideas about the aspect that differs.

Dissonance is the term used for this divergence from a well-understood concept into an alien one. The Tharps describe it as a "useful friction generated through the act of interpretation somewhere between the truly strange and the familiar" (403). The familiarity of most of the scenario enables thoughtful engagement; the part that differs is the intended source of discourse.

This style of object creation can be thought of as a partnership between designer and audience member, who is not necessarily the prospective user. In fact, in many cases, including this project, the user is a means to discussion than anything else:

> Users are usually more actor-like, more propositional, more fictional; they are conjured up as a way to better understand the artifact and the scenario. Primarily, discursive designers seek to 'sell' ideas, not products. To help mitigate this common confusion, we refer to the actor in the scenario who uses that artifact as a 'rhetorical user', in the sense that they are a rhetorical device of sorts. (189)

The audience, then, is a separate entity to the user in the formulation of discursive design. For this project, the audience likely includes future users of the product being discussed. Many people will go on to own or use autonomous vehicles one day; it is this fact that makes challenging them as a contemporary audience so worthy.

One more explicative function of discursive design is its ability to cut through the weeds of information overload and pick out a stark, compelling example. As the Tharps write:

> One of the challenges of the information age is that there is access to more information but not any more cognitive ability to deal with it. Design can, however, offer clear and cogent ways to package and communicate new, relevant, and interesting information for consumption...Discursive design swims against the tide and is asking for more attention and work from the audience. (159)

This mode of design suits this project well because its goal is audience reflection and circumspection. While a commercial design project on the same subject might propose a new kind of helmet or safety mechanism for pedestrians sharing roads with autonomous vehicles, it would miss an essential element of having the discussion today: it affords the audience time and space to consider their relationship to these scenarios, and has them question whether they consent to have them emerge as depicted.

Grounded Theory and Zinemaking as a Method of Practice

While discursive design offers a framework through which to evaluate the produced work on an ongoing basis, it does not offer a methodology for the production of this work. In consultation with my advisors, the work to sharpen the scope of this project was primarily done through iterative zinemaking. This zinemaking, in combination with a version of grounded design principles, allowed for the emergence of unifying themes and modes of inquiry around which specific outcomes could coalesce.

Grounded design is an inductive process allowing the researcher to produce hypotheses as part of collating and "coding" qualitative data that may initially seem to lack patterns (Glaser). By taking subsequent passes through freeform data, the researcher slowly abstracts further and further from the data itself into a series of conjectures. Although this process is much more rigorous in conventional application, such as information technology research (Weatherall) and nursing research (Groves), the concept of assembling design approaches from sets of freeform notes and ideas was helpful.

First, I laid out terms, concepts, threats and opportunities on paper, seeking to identify the most accessible way to draw parallels between seemingly disparate notions. From there, I selected what I thought were the most successful topics of discussion and started to create simple paper zines that demonstrated one approach to stimulating a conversation about these matters among the public.

This approach helped me envision a speculative built world. Imagining marketing approaches for the integration of autonomous personal appliances into everyday life helped me to construct a hierarchy of what elements people were most likely to find endearing or approachable about these appliances, and what characteristics would engender the most resistance.

The appliances I designed covered a variety of functions, and several more designs were proposed but never turned into zines. The product repair manual was an example of an effective metaphor to help embed these elements as functional, believable and fallible in the future world of autonomy.

The zines that I felt hinted at the most complex set of follow-up questions were ones that I chose to redraw digitally at higher resolution. This technique of increasing fidelity while keeping content similar was a way to isolate and evaluate the quality of the idea as a discursive object. Some elements that were added to the digital zines became outcomes in their own right. For example, the fiduciary calibration markers in the corners of road signs (see Appendix B) were intended to show how an autonomous vehicle's vision module might lock onto and parse a sign such as this, but in designing this sign it became obvious that the applications of these markers were wider and deserving of their own discussion.

In this way the zinemaking process is mitogenic, spawning concepts worthy of their own discussion in the process of discussing other ideas. Similarly, the concept of arranging road access times according to vehicles' emissions impact and mode of control was inspired by London's Congestion Charge and Ultra Low Emissions Zone programs (Transport for London), but eventually

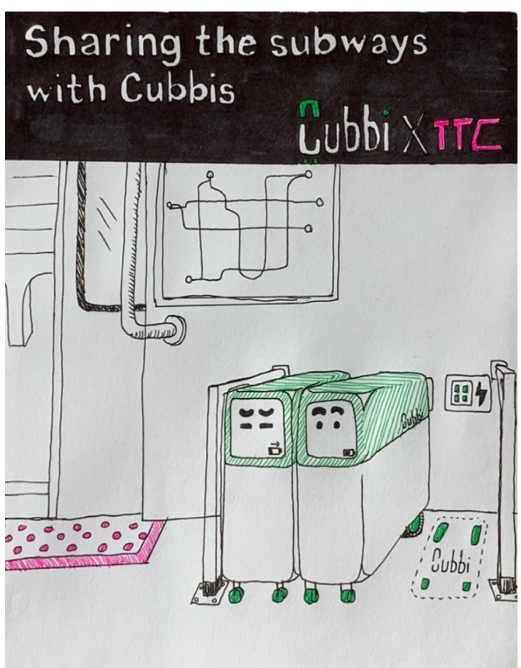

Figure 4 Early zines focused on integration of autonomous appliances with daily urban life. Special attention was paid to anthropomorphization tactics often used to build trust and comfort in prospective users.

blossomed into a wider exploration of the gradations of automotive autonomy and how vehicle capabilities and responsibilities at these levels are regulated.

Diorama and Plasticity as Models of Viewership

Early in the ideation process I landed on the concept of an exposition as an encapsulating metaphor for the project. Expos and World's Fairs have long been vehicles for futuristic ideas, using their grand scale and ambition to whisk audiences away to fantastical worlds, often superimposed on images of the very city they are in.

Although the concept of expo as encapsulation was rejected during the iteration process, the core idea remained: how do you sufficiently remove an audience from their current circumstances to achieve the plasticity needed to take on such foreign ideas and concepts? Expos do this with grandeur and sleight of hand: they transform industrial wasteland into gleaming metropolis, if only for a few months. Vanishingly few exposition grounds manage to incorporate their boondoggles into the fabric of their cities twenty years later, and fewer still deliver on the promises of generational economic and cultural revitalization made to citizens during the pitch but, for the duration of the exposition, the audience is transported elsewhere.

This project has access to far fewer resources, but I searched for ways to lift something from the exposition modality to better immerse the audience in the dense, foreign ideas that I was trying to convey. Reviewing Norman Bel Geddes' successes as creator of dioramas was key to understanding what captivates audiences about futurism: believability and scope. Bel Geddes had a fastidious commitment to rendering the diorama as a replica of existing American society. "[He] and his staff meticulously studied the photographs to establish the scale and environmental effects of various geographic and urban elements as seen from an airplane" (Morshed 75).

In basing my outcomes in actual Toronto intersections, especially the interactive outcomes that are situated in the intersections near the OCAD University campus, I aim to bootstrap the *lived-in-ness* of the city to better enhance the credibility of the designs themselves. The intersection is the discursive scenario, and the autonomous vehicles provide the discursive friction.

I believe that superimposing these new digital elements over an existing urban scene functions as a kind of contemporary diorama, providing in the viewer a distancing effect that allows them to appraise the implications of these changes to the spaces they navigate through daily.

This distancing must occur in two dimensions for the best plasticity: time and scope. For an audience to latch onto and be willing to engage with a speculative idea, the scenario described must fulfil a tricky balancing act: not so far away as to be hard to envision, but not so soon as to be infeasible. A viewer must be able to intuitively trace the steps from current day to the speculative scenario. Similarly, in terms of scope a designer can produce a scenario either too small or too big. Too small, and they threaten to bore or patronize the viewer; too big and they risk inducing either panic or apathy. We see this often with scenarios that envision climate change or nuclear war; the scale of change is so all-encompassing that critique becomes useless. An audience must feel empowered to aim towards or away from the scenario provided within the timeframe described.

This framework and the diagram below (see Figure 5) take their inspiration from a futures cone, a diagram originally developed in the 1990s by Charles Taylor, Trevor Hancock and Clement Bezold ("Discursive" 200). This expanding volume demonstrates the escalating possibilities and plausibilities of speculative fiction as the time distance between the scenario and the present day grows. What the diagram below does is effectively take a two-dimensional section of this cone and restrict it further along the temporal axis. It is not enough to keep the scenario plausible; to achieve audience plasticity requires that a discursive scenario be placed a specific, measured amount of time into the future.

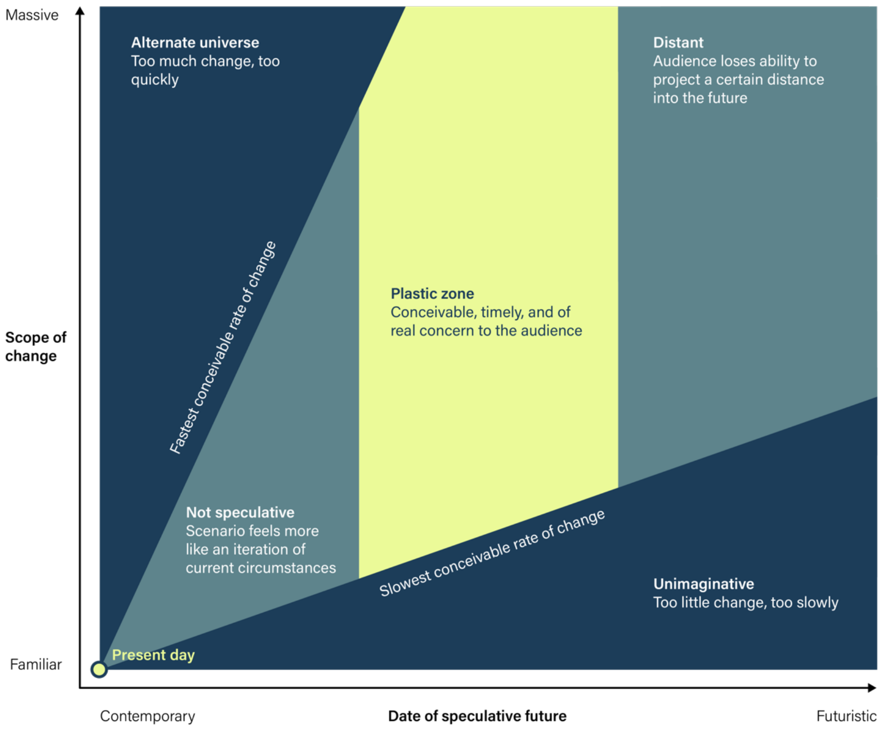

Figure 5 A model for envisioning speculative scenarios that induce maximum plasticity, or openness, in the viewer.

This project aims to show the viewer a world twenty or thirty years away which overlays and augments their existing world. Grounded and distanced, the viewer isolates the aspects of the scenario that would not affect their life and can more readily appraise the aspects that would. For example, the interactive component of the project uses an existing, popular intersection as its home base, at the corner of Richmond and Duncan Streets in Toronto. The area includes a popular cycle path, several lanes of traffic, and traffic signals. At least four modes of transit intersect in this square of land, allowing the viewer to engage with the potential impact on all of them. Far from being some featureless simulated environment, many viewers have memories, commutes, journeys and stories associated with this area, allowing them to project the changes I envision both forwards and backwards in time.

Project

The design outcomes to this project comprise a series of posters, each depicting and exploring a different technical, ethical, or physical dimension of the discussion surrounding autonomous vehicles, and a set of interactive webpages designed to be experienced either in the gallery, in the intersection directly outside the gallery, or remotely on any smartphone or computer.

Refining Outcomes

With discursive design as a guiding methodology and plasticity and diorama as modes of viewership, the key to selecting outcome formats became finding self-effacing media that would allow the content and not the form to take precedence in discussion. The outcomes needed also to negotiate the space between the present day and the near future, and to remain grounded in the local built environment.

Using a simplified form of grounded theory, I began by mind-mapping concepts and concerns from my research and collecting them into like groups (see Appendix A). Those groups became low-fidelity zines, which were critiqued and evaluated for discursive potential (see Appendix B). A second round of zinemaking involved regrouping concepts and splitting others into multiple zines (see Appendix C).

Initial presentations were made to faculty and advisors, along with industry subject matter experts at the Canadian Film Centre's Media Lab. At this stage, the encapsulating metaphor had emerged as a speculative world exposition, hosted in Toronto in

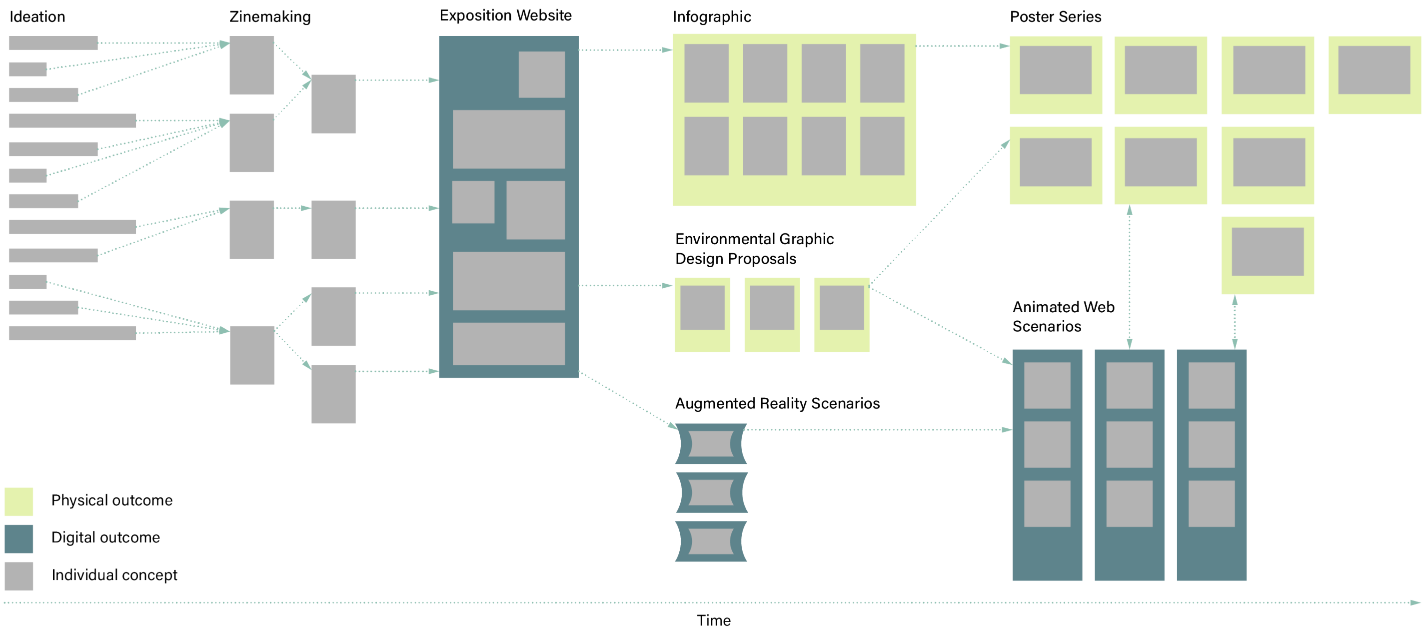

Figure 6 Over time, the format of the design responses to the project critique evolved. In the final format, outcomes dialogue across multiple media, producing a cohesive response in both physical and digital format.

the year 2067 (see Appendix D), a century after the city's last expo. I believed that this capsule provided the distancing effect I wanted, while grounding the viewer in a city known to them. It became clear in discussion that, though the exposition was a well-known event format on which speculative ideas could be projected, the concept had two main drawbacks: the exposition as a historical event comes with its own baggage and implications worthy of discussion in their own right, and getting the desired degree of verisimilitude in an exposition website was likely beyond my programming ability, to the detriment of the end product.

I made the decision after that to focus on two modes of output: physical and augmented reality. Physical outcomes would provide a tangible, dense discussion of the issues at heart while also functioning well in exhibition space. Augmented reality would provide a futuristic and visceral experience to complement the physical.

In developing a single, overarching infographic for display in gallery space, the quantity of content that emerged for discussion necessitated its splitting into several posters. This division had narrative consequences too: the poster series needed to involve a through-line that creates an effective narrative, and the question of what information should be disseminated and in what order became paramount.

At the same time, development work was ongoing on the augmented reality experiences that would complement the poster series, but the zero-dependency, widely accessible, browser-only version I wanted to develop had significant performance issues on most devices, triggering a shift to a more conventional webpage design. This design attempts to combine the tactility of the augmented reality viewing experience with additional stability and interactive control.

The final outcomes combine physical and digital formats (see Figure 6), focusing on delivering a multi-modal experience that attempts to maximize the viewer's engagement with each individual topic of discussion.

Exhibiting the Project

As a discursive project, this project's effectiveness hinges on its capability to provoke awareness in the viewer. In Tharp's model, this project is *instrumental-external*; rather than existing as a design object in and of itself, it is an instrument through which the external audience can construct their own critique. The project must present the principles of autonomous driving proliferation in a self-effacing way, nudging the viewer towards the intended conclusion rather than putting an entire curriculum on rails.

In response to the context of the problem, the project's design outcomes use several visual strategies across a range of media. This multi-pronged approach mimics the polyglot expectations of the future: as navigators, we will need to process and respond to wildly varying inputs and data views. In some sense, one project delivered multiple ways is a good approximation of the cognitive task of moving around a city in twenty years' time.

In the exhibition space, twelve posters, twenty-four by eighteen inches each, are mounted in two rows of six (see Figure 7). A projector throws a looped video playthrough of the two web experiences into an identical twenty-four by eighteen inch footprint to the posters' right. Above the projection sits a poster displaying QR codes and URLs for viewers to open and use the web experiences themselves, along with contact details for the artist. This interplay of media simulates the experience of divided attention at intersections, where multiple streams of input data must be processed by the human brain to ascertain whether crossing is safe.

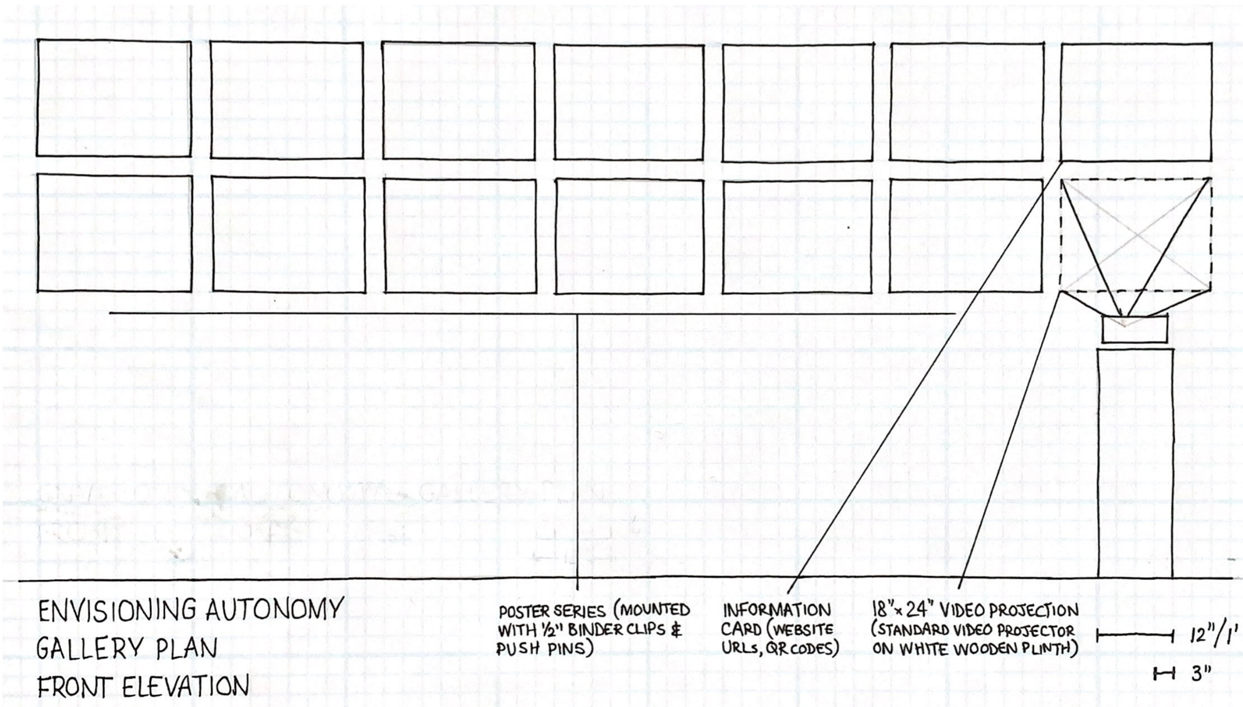

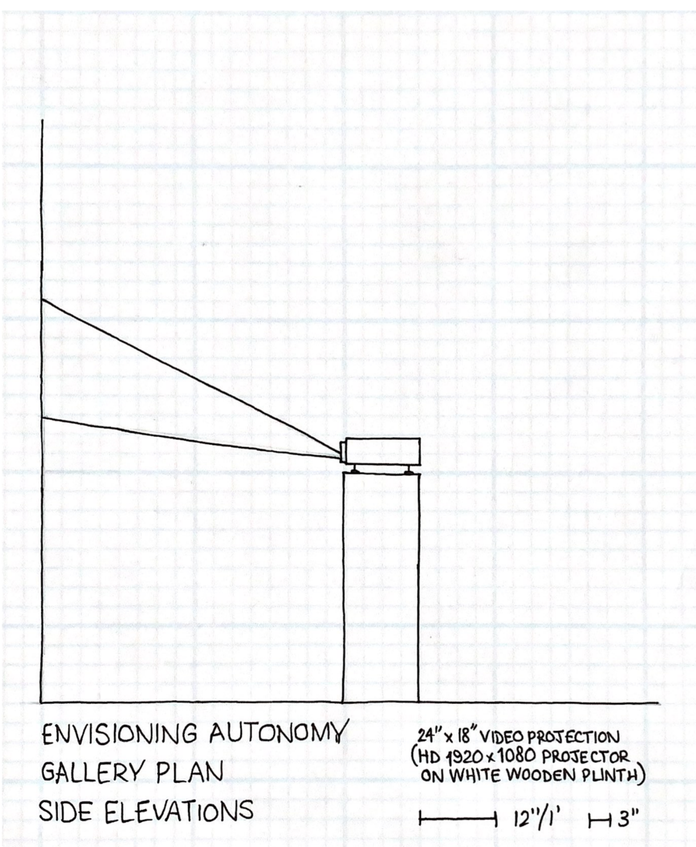

Figure 7 A gallery plan that allows the exhibition of the poster series and a looped web animation simultaneously. The interplay of media streams mimics the division of attention experienced at a crosswalk.

Figure 8 (next page) Images from the project's initial exhibition show its scale and placement in physical space.


Envisioning Autonomy
What Makes a Road?
Pedestrian Signal Head for Autonomous Vehicles
Levels of Autonomous Driving
How Do AVs Sense the World?
Augmented Reality Visualizations for Pedestrian Safety


Pedestrian Signal Head for Autonomous Vehicles
Intersection Management
Would My Car Save the Same People As Yours?
Ongoing Research Approaches
Augmented Reality Visualizations for Pedestrian Safety
Intersection Management
The Moral Machine
What Happens Next?

Poster Series

Format Evolution

The poster series emerged as the conclusion of the zinemaking process. Throughout the early months of research into autonomous vehicles and the infrastructure behind their operation, a set of topics emerged that felt like a strong general survey of the research space. Once reviewed, the viewer could conceivably have an informed conversation with a proponent of the technology.

Although early zines focused on autonomous appliances and the gradual equipment of household gadgets with arrays of sensors and cameras, the speed and potential of autonomous driving quickly emerged as the most urgent topic of focus.

Initial designs focused on sections of a single, large poster output, but it became apparent that there was too much information for a single poster. With the additional breathing room provided by a poster series, questions of narrative, information density, and tone positioning could be more fully considered. With individual elements constructed, a manual paste-up process helped to visualize the appearance of each poster in situ (see Figure 8).

With so many posters conveying such a large amount of information, chief considerations became density, flow and narrative. What order of dissemination would develop the most plastic frame of mind in the audience? In environmental graphic design, this same concept is foundational to its aims: progressive disclosure. Much research has been done into the nature of how humans learn both individually and as a collective. It is important not to overwhelm the viewer with overly dense compositions, but to thoughtfully add depth and nuance over time. Of course, in many examples of graphic design the intent is to overwhelm the viewer, but the specific aims of this project demand an engaged, unhurried audience familiarizing themselves with the subject matter.

Other considerations are paramount for an exhibited visual project such as this one. Viewing distances, viewing heights, and the natural movement of the eye are all potential hindrances to the optimal viewing experience. The gallery plan for this exhibition attempts to account for all these concerns.

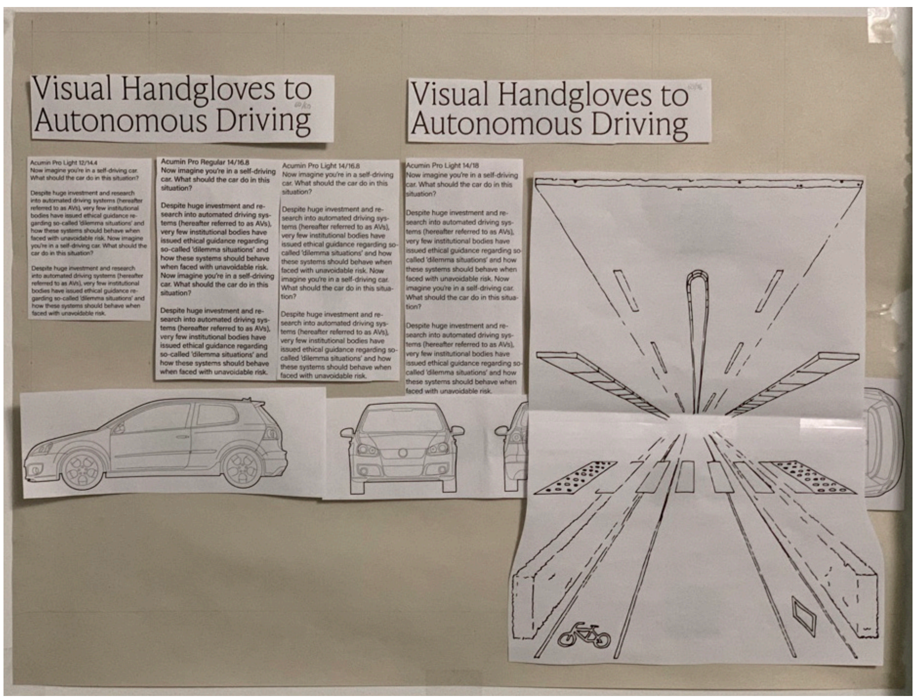

Figure 9 During paste-up, individual elements are composed at scale, evaluated, and adjusted for visual acuity and balance. Test phrases like "Handgloves" are used to ensure ascending and descending characters do not interfere in context.

Discursive Intent

The goal of this poster series is not only to inform but to situate viewers in a mindset of transition. The discursive design approach allows these transitional states – the migration towards autonomous driving, the gradual addition of Internet of Things-connected sensors to our daily lives, and intersections and crosswalks themselves – to be captured and expanded for analysis. Phases of movement that are usually taken as read or elicit little reaction can be thoroughly examined in this format.

The series also includes elements of environmental graphic design; this is the closest the project comes to suggesting physical interventions in the built environment as strategies to counter some of the concerns identified. A proposal for an augmentative signal head, to be used by pedestrians at crosswalks to predict movement patterns of sidewalk-using autonomous appliances, includes design intent drawings that demonstrate how the equipment might function at scale.

The goal of these elements of environmental graphic design is to demonstrate the real possibilities that physical signage and lighting offer in terms of protecting and entrenching a balance of access to safe navigation for pedestrians. These elements act as bulwarks against the progressive encroachment of cars and other motorized vehicles on roads and streets.

A strip of unmarked tarmac	Provides a flat, consistent surface. Makes travel possible for consumer vehicles. Reduces tire degradation.
A raised median	Separates oncoming traffic flows, preventing head-on collisions.
Painted lanes	Controls lateral movement within traffic flows. Provides visual guides to help drivers position vehicles.
Curbs and sidewalks	Sets outer boundaries of the road. Physically separates pedestrians from the flow of traffic.
Signage	Contextualizes location and conveys laws related to speed, parking, and occupancy limits.
Crosswalks and traffic lights	Mediates intersecting travel, enabling other vehicles to cross the road safely. Signals mutually agreeable zones of interchange.
Hedges and greenery	Mitigates noise, smell and air quality caused by road traffic.
Cyclist and protected lanes	Protects the motion of certain types of traveller, often for safety or to incentivize low-carbon travel.

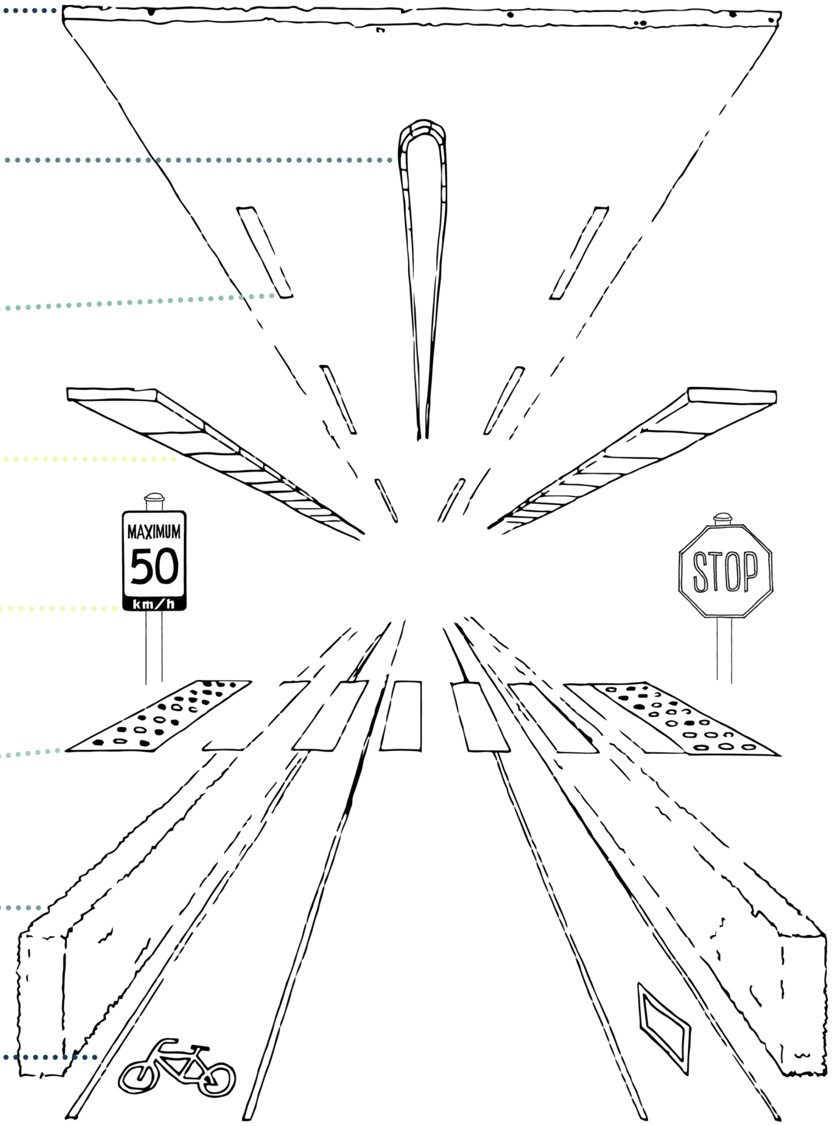

Figure 10 A portion of an individual poster from the series. This sheet asks the viewer to consider what functions are achieved by individual elements of road design, and how autonomous vehicles might respond to each of these.

Outcome Description

Initial panels explore the levels of autonomy afforded by certain combinations of onboard sensors and computers, from contemporary driver assistance features (Level 1) all the way up to fully autonomous operation with no human steering or motion inputs (Level 5). The sets of sensors that facilitate autonomous driving are described individually, demonstrating that these vehicles are processing information across a wider subset of the sound-light spectrum than even humans do. To end the first group of panels, roads themselves are split out into their constituent elements and typified by function, to better understand the composability and flexibility of road design for correct use cases.

Snapshot summaries of dynamic project outcomes follow: augmented reality strategies for highlighting high-momentum vehicles and other pedestrian dangers, a pedestrian signal head for use by sidewalk-occupying autonomous vehicles, and various intersection phase designs for use in full autonomy scenarios.

Finally, the series reviews the ethical concerns of algorithm-driven decision making in life and death scenarios, including recent studies showing global sentiment towards such decision making, before reviewing existing research approaches in the field to resolving these issues.

The goal of the poster series is to lead the viewer to a point of circumspection. Although the quantity of instructive material may be quite high, it aims to strike a fine balance of progressive disclosure so as not to merely overwhelm. While engulfing the viewer in moral dilemmas and technical minutiae would have the somewhat desired effect of inducing caution, what the project seeks to endow instead is an informed, critical stance.

Web Experiences

Contextual Applications of Augmented Reality

Complementing the poster series is a collection of three interactive web experiences that aim to render the dynamic elements of the project in a more immediate way for viewers.

These webpages have their root in augmented reality experiences. What the project requires is a technology that allows viewers to envision their own daily lives modified by the addition of autonomous vehicles. Augmented reality fits this use case very well: it is an innately immersive, supplementary and futuristic technology. As an emergent medium, it provides both familiarity and unfamiliarity to the viewer, showcasing its potential to maximize audience plasticity to near speculative futures. Much like autonomous driving technology, augmented reality asks us to renegotiate our relationships to existing spaces and the processes we use to get around them.

Augmented reality is destined to become our main mode of engagement with the Internet, as wearable devices become more accessible and powerful (Porter 104), and immersive experiences become the benchmark in gaming, entertainment and marketing.

Part of the appeal of functional augmented reality is its reduction of cognitive load, another concept borrowed from the world of wayfinding. By attaching floating data points to real-world elements and entrenching them within real space, the technology reduces the mental distance that the viewer must travel to read and understand virtual concepts and attach them to tangible outcomes. In this sense, the technology functions as an element of environmental graphic design, attaching meaning to real space and making its significance explicit.

In higher-end vehicles, technologies adjacent to AR are already available. “Heads-up” windshield displays overlay the real world with driving data like speed and upcoming navigation tasks, improving safety by reducing the cognitive load of the driving task and increasing the amount of time the driver spends with their eyes on the road ahead. Smart wingmirrors also offer contextual information to drivers like blind spot warnings and overtaking cameras, further enhancing and augmenting the driver's ability to picture hazards and obstacles around them.

Using Conventional Web Technologies

It became clear during the development of prototypes for augmented reality web experiences that zero-dependency, browser-based AR is still in its infancy, with poor marker recognition causing such significant jitter of onscreen assets that reading and processing the information was all but impossible. Based on the principle of the project that informational delivery methods should be self-effacing and foreground the information itself, a pivot to more conventional web technologies was necessary.

With the more fully developed and stable technologies available in established web development, an approximation of augmented reality can be designed to effectively simulate its eventual seamlessness and detail, rather than approximate it with current technologies (see Figure 10). By prioritizing the approximation of a certain visual goal – the quality of the overlaid elements blending seamlessly into the built environment – over precise fidelity to the technical of augmented reality, the project can communicate its concerns and critiques more effectively.

As Bruce and Stephanie Tharp write in *Discursive Design,* "while it may be nice to have a fully functional prototype, should the resources needed for crafting the artifact take precedence over crafting the discourse?...Fundamentally, it is important for designers to understand the challenges and opportunities involved and that adjustments may be necessary regarding their intentions, approaches, and concerns for efficiency and risk ("Discursive" 168-9).

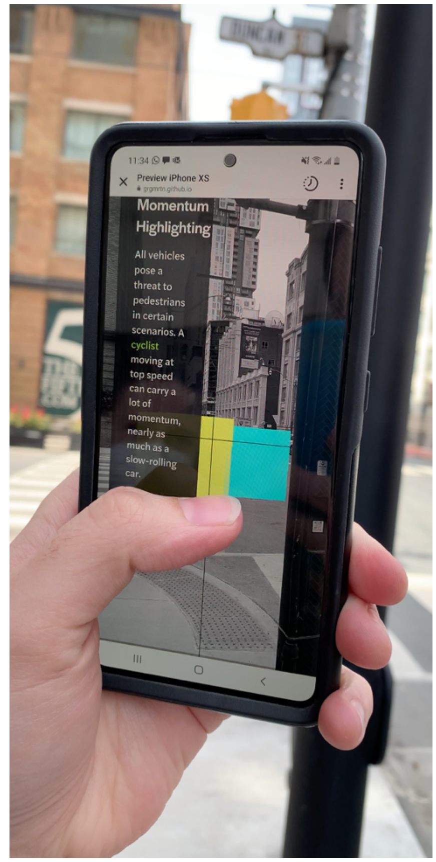

Figure 11 The webpages are designed to work both in the built environment on smartphone displays and in the more conventional desktop environment.

Outcome Description

Using conventional and lightweight web technologies, these webpages render animated scenarios related to autonomous driving controlled entirely by the user's scroll position on page. Using this input technique, the viewer maintains ultimate control over the flow and direction of information. The THREE.js library is used to generate three-dimensional assets that conform to the perspective of the background image. These background images depict locations near the exhibition space and are referenced using QR codes on nearby posters.

The positioning of these posters is designed to produce a matching perspective on the user's device that allows them to envision the screen itself as a portal into this speculative future. The QR codes are placed at the requisite height to produce as close a perspective match as possible to the images onscreen (see Figure 11).

Figure 12 Posters in city intersections are designed to look like actual provincial notices, to improve their verisimilitude and entice the viewer to access the webpage.

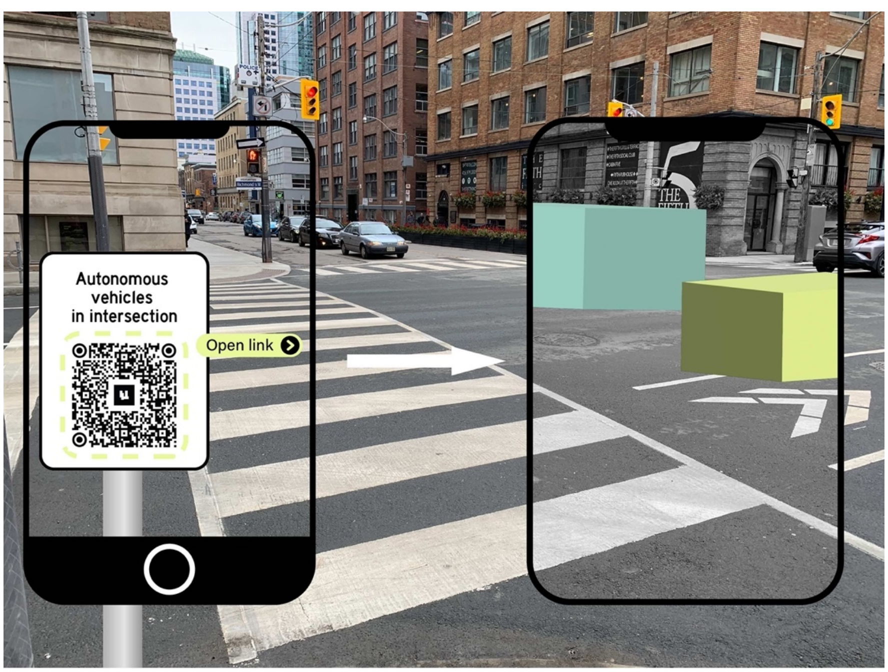

Figure 13 The interactive webpage components are activated by a printed QR code, placed at the intersection to mimic the perspective of the static black and white background.

Animated elements overlay the background, allowing the viewer to better envision autonomous infrastructure in their daily lives.

Dialogues between Design Outcomes

These webpages are designed to exist in conversation with the poster series. Two of the three scenarios depicted in the webpages are replicated in a modified format for the poster series. This allows additional retention of the concept by the audience and mitigates the weaknesses of each visual format. What the poster series might lack in its ability to produce a visceral reaction is countered by the webpages' immediacy and capability for fluid interaction. What the webpages may lack in visual fidelity or explanatory detail is countered by the poster series' additional density and more didactic tone.

By using a coherent visual system of type, colour and iconography, these outcomes are pulled even closer together into a multi-modal discursive object. Much as the visual interventions necessary for the safe rollout of autonomous driving necessitate a variety of modes for their optimal function, the design objects produced by the project require several modes to communicate their concerns effectively.

Figure 14 In desktop environments, the additional screen real estate is used to provide additional detail and an easier viewing experience.

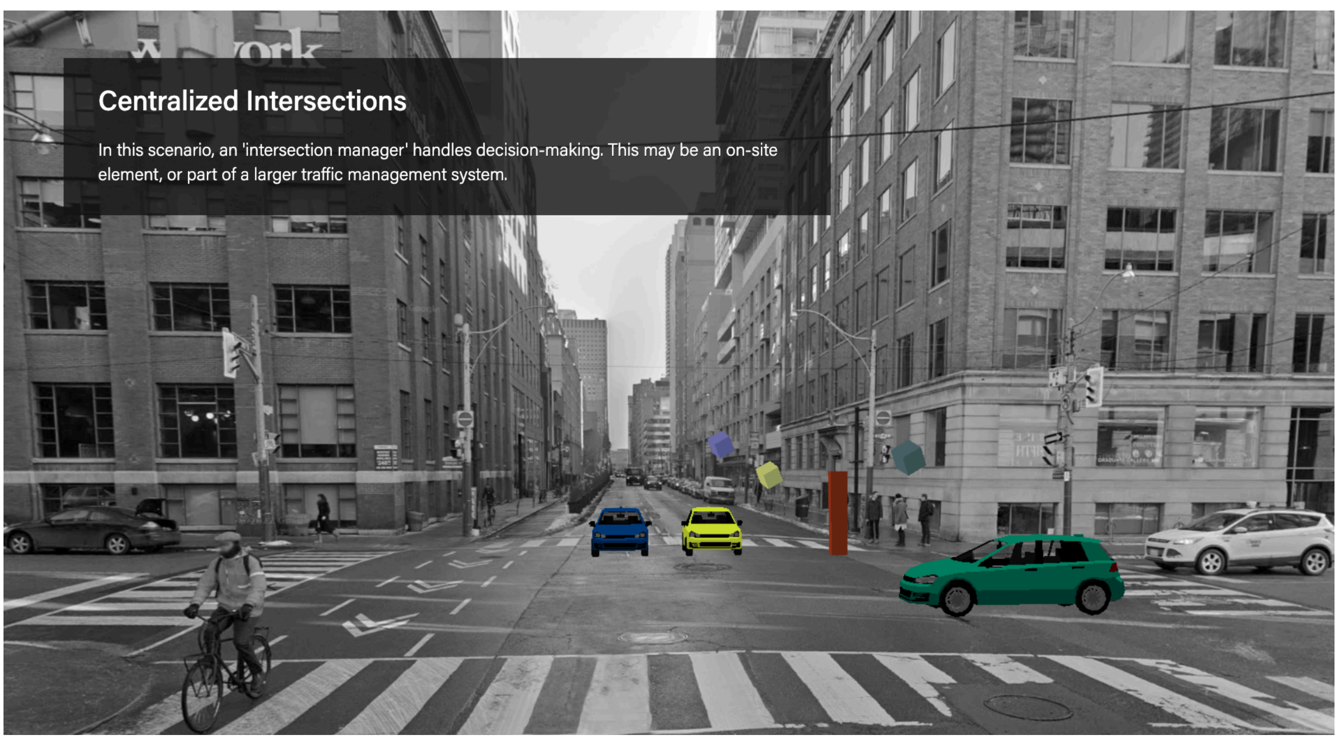

Conclusion and Future Work

Autonomous driving will revolutionize many aspects of human life. Private car ownership, the trucking industry, and the nature of commuting all stand to be upended by the changes this technology will bring. What I have argued above, and what the design outcomes have attempted to illuminate, goes further: the fundamental act of mobility itself is changing.

Whether or not you own an autonomous car, the changes we make to our roads and cities to accommodate their proliferation will have a significant impact. Crosswalks, where pedestrians may already feel at a disadvantage, become nexuses of indecipherable, affectless vehicles, moving less predictably than ever as they dispense with the visual vernacular of waves and eye contact that many of us use to get around.

The insertion of artificial intelligence into all aspects of route planning has implications that modern navigation apps have not yet asked us to reckon with. If it transpires that all autonomous vehicles are receiving instructions from a centralized management entity, a serious risk emerges that certain streets, intersections, and entire neighbourhoods can be systematically avoided or starved of traffic and footfall.

The race to get where we're going as swiftly as possible has always carried risk. Constant attention can occasionally waver, and this can lead to serious accidents, some where the driver is at fault and some where pedestrians and passengers become unfortunate victims of chance.

Now, with autonomous vehicles, we will transition to a world where cars pay constant, minute attention to every object nearby and can react quickly to unexpected changes in the environment. Rather than consign fatal accidents to history, however, these vehicles will make the few accidents that remain even more fraught. Assigning responsibility, maintaining safety standards and developing agreeable ethical principles becomes more difficult in a world of massive datasets and indecipherable algorithms.

What this project has tried to show is that, though the technology may be inevitable, the nature of its implementation is still very much up for debate. Many of those with influence in this dispute are well-positioned to enrich themselves from the rise of autonomous vehicles. We cannot rely on crises of conscience to deliver the transparent oversight we need.

What tools do we have at our disposal in the coming fight to educate and regulate? Many of the details are beyond the scope of this project, but at a minimum, we must expect more technical literacy and effective oversight from our governments. We must learn from multiple recent examples of the corporate world running rings around those who ostensibly regulate their activity.

Public awareness is the most powerful tool at our disposal. Without concerted information campaigns, many people's first encounter with no-win driving scenarios will be their participation in one. The gradual scrubbing from view of raw data and interface friction may leave us dependent on machines we don't understand, own or control.

While this project has focused on the specifics of visual communication in the urban environment, interventions for the visually impaired become paramount in a world where machines may no longer easily recognize those in need of additional courtesy and notice. Auditory and tactile interventions play an important role in making our streets safe and accessible for all.

Next time you wait at a crosswalk, concentrate on the ways we communicate with each other that have nothing to do with the machines we are operating. These makeshift dialects are vital components of urban lifestyle, and we must protect them as we should any other endangered language.

Works Cited

"Roads: 2021 Infrastructure Report Card." *American Society of Civil Engineers,* 2021. https://infrastructurereportcard.org/wp-content/uploads/2017/01/Roads-2021.pdf. Accessed 9 December 2021.

"AutoGrasp™ Robotics Intelligence Platform." *Kindred AutoGrasp,* https://www.kindred.ai/autograsp. Accessed 9 December 2021.

"CO2 Emissions (Metric Tons per Capita)." *World Bank DataBank*, 2020. https://data.worldbank.org/indicator/EN.ATM.CO2E.PC. Accessed 28 November 2021.

"Fatalities – Vision Zero." *City of Toronto*, 18 November 2021. https://www.toronto.ca/services-payments/streets-parking-transportation/road-safety/vision-zero/vision-zero-dashboard/fatalities-vision-zero/. Accessed 21 November 2021.

"GPS Accuracy". *GPS.gov*, 6 August 2021. https://www.gps.gov/systems/gps/performance/accuracy/. Accessed 26 November 2021.

"Signatories & Founding Coalition." *Shared Mobility Principles for Livable Cities*, 2017, https://www.sharedmobilityprinciples.org/home/#signatories. Accessed 9 December 2021.

"Strong AI." *IBM Cloud Education*, 31 August 2020. https://www.ibm.com/cloud/learn/strong-ai. Accessed 21 November 2021.

"The Boring Company | Tunnels." *YouTube,* uploaded by The Boring Company, 28 April 2017. https://www.youtube.com/watch?v=u5V_VzRrSBI. Accessed 9 December 2021.

"Vision Zero 2.0 – Road Safety Plan Update." *City of Toronto*, 13 June 2019. https://www.toronto.ca/legdocs/mmis/2019/ie/bgrd/backgroundfile-134964.pdf. Accessed 26 November 2021.

"Why Principle 10." *Shared Mobility Principles for Livable Cities*, https://www.sharedmobilityprinciples.org/why-10. Accessed 9 December 2021.

Adams, Douglas. "Norman Bel Geddes and Streamlined Spaces." *Journal of Architectural Education*, 30:1, 22-24. 1976. https://doi.org/10.1080/10464883.1976.10758072.

Barter, Paul. "'Cars are Parked 95% of the Time.' Let's Check!" *Reinventing Parking*, 22 February 2013. https://www.reinventingparking.org/2013/02/cars-are-parked-95-of-time-lets-check.html. Accessed 19 November 2021.

Bastani, Aaron. *Fully Automated Luxury Communism: A Manifesto*. Verso, 2020.

Beresford, Colin. “Pedestrian Deaths in 2019 Were Highest in 30 Years, Report Says.” *Car and Driver*, 27 February 2020. https://www.caranddriver.com/news/a31136893/pedestrian-deaths-increase-2019/. Accessed 21 November 2021.

Boudette, Neal E., and Niraj Chokshi. “U.S. Will Investigate Tesla's Autopilot System over Crashes with Emergency Vehicles.” *The New York Times,* The New York Times, 16 Aug. 2021, https://www.nytimes.com/2021/08/16/business/tesla-autopilot-nhtsa.html.

Brodkin, Jon. “Ajit Pai on Net Neutrality: ‘I favor an open Internet and I oppose Title II’”. *ArsTechnica*, 31 January 2017. https://arstechnica.com/tech-policy/2017/01/fcc-chair-ajit-pai-wont-say-whether-hell-enforce-net-neutrality-rules/. Accessed 28 November 2021.

Brueninger, Kevin. “Biden Signs Order to Crack Down on Big Tech, Boost Competition 'Across the Board'.” *CNBC Tech*, CNBC, 9 July 2021, https://www.cnbc.com/2021/07/09/biden-to-sign-executive-order-aimed-at-cracking-down-on-big-tech-business-practices.html. Accessed 28 November 2021.

Chappell, Robert and Bringhurst, Warren. *A Short History of the Printed Word*. Hartley & Marks, 1999.

Coll, Steve. *Private Empire: ExxonMobil and American Power.* Penguin, 2012.

DiMento, Joseph F. C., and Cliff Ellis. *Changing Lanes: Visions and Histories of Urban Freeways.* MIT Press, 2014.

Efrati, Amir. “Money Pit: Self-Driving Cars' $16 Billion Cash Burn.” *The Information*, The Information, 21 Dec. 2020, https://www.theinformation.com/articles/money-pit-self-driving-cars-16-billion-cash-burn.

Eyre, Nick. *Shifting the Focus: Energy Demand in a Net-Zero Carbon UK.* The Centre for Research into Energy Demand Solutions (CREDS), 2019.

Glaser, Barney G. *The Discovery of Grounded Theory: Strategies for Qualitative Research.* Aldine, 1967.

Goldberger, Paul. “Back to the Future: A New Look at Modernist Hero Norman Bel Geddes, Designer of the Original 1939 ‘Futurama.’” *Vanity Fair*, 22 Oct. 2013, https://www.vanityfair.com/culture/architecture/2013/10/norman-bel-geddes-designer-original-futurama. Accessed 9 December 2021.

Goldhill, Olivia. “We've Had Driverless Cars for Almost a Hundred Years.” *Quartz*, 22 Oct. 2016, https://qz.com/814019/driverless-cars-are-100-years-old/. Accessed 24 February 2022.

Graham, Stephen. *Vertical: the City from Satellites to Bunkers*. Verso, 2016.

Greenfield, Adam. *Radical Technologies: the Design of Everyday Life*. Verso, 2018.

Groves, Patricia S et al. "Handing Off Safety at the Bedside." *Clinical nursing research vol.* 25, no. 5, 2016: 473-93. https://doi.org/10.1177/1054773816630535.

Heinrichs, Dirk. "Autonomous Driving and Urban Land Use." *Autonomous Driving*, 2016, pp. 213–231., https://doi.org/10.1007/978-3-662-48847-8_11.

Jacobs, Jane. *The Death and Life of Great American Cities*. Vintage Books, 1961.

Kopytoff, Verne. "Are Google views too revealing?" *SFGate*, 2 June 2007. https://www.sfgate.com/business/article/Are-Google-views-too-revealing-Online-maps-2590209.php. Accessed 19 November 2021.

Kurzweil, Ray. *The Singularity Is Near: When Humans Transcend Biology*. Viking, 2005.

Lynch, Kevin. *The Image of the City*. MIT Press, 1960.

Mariani, Stefano, et al. "Coordination of Autonomous Vehicles." *ACM Computing Surveys,* vol. 54, no. 1, 2021, pp. 1–33., https://doi.org/10.1145/3431231.

Matz, Carlyn J et al. "Evaluation of daily time spent in transportation and traffic-influenced microenvironments by urban Canadians." *Air quality, atmosphere, & health*, vol. 11, 2 (2018): 209-220. doi:10.1007/s11869-017-0532-6

Morshed, Adnan. "The Aesthetics of Ascension in Norman Bel Geddes's Futurama." *Journal of the Society of Architectural Historians,* vol. 63, no. 1, [Society of Architectural Historians, University of California Press], 2004, pp. 74–99, https://doi.org/10.2307/4127993.

Passini, Romedi and Arthur, Paul. *Wayfinding: People, Signs and Architecture*. 1992.

Porter, Michael E., and James E. Heppelmann. "Why Every Organization Needs an Augmented Reality Strategy." *Harvard Business Review's 10 MUST:* 85-104.

Rifkin, Jeremy. *The End of Work: the Decline of the Global Labor Force and the Dawn of the Post-Market Era*. Tarcher/Penguin, 2004.

Shepardson, David. "Tesla Recalls Nearly 54,000 Vehicles That May Disobey Stop Signs." *Reuters*, Thomson Reuters, 1 Feb. 2022, https://www.reuters.com/business/autos-transportation/tesla-recalls-nearly-54000-us-vehicles-rolling-stop-software-feature-2022-02-01/.

Silva, Dahlen et al. *Autonomous Vehicles and Urban Space Management*. Silesian University of Technology, 2021. https://www.researchgate.net/publication/348631287_Autonomous_Vehicles_and_Urban_Space_Management

Singh, Santosh. "Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey." *Traffic Safety Facts CrashStats*, vol.

812, no. 115, Feb. 2015, https://crashstats.nhtsa.dot.gov/Api/Public/Publication/812115. Accessed 22 Feb. 2022.

Steinbeck, John. *The Grapes of Wrath*. Penguin, 2008.

Stromberg, Joseph. “Highways Gutted American Cities. So Why Did They Build Them?” *Vox*, Vox, 14 May 2015, https://www.vox.com/2015/5/14/8605917/highways-interstate-cities-history.

Sutter, John D. “Income Inequality and Google Maps.” *CNN Opinion,* 7 November 2013. https://www.cnn.com/2013/11/07/opinion/sutter-google-maps-income-inequality/index.html. Accessed 19 November 2021.

Tharp, Bruce M., and Stephanie M. Tharp. “Discursive Design Basics: Mode and Audience.” *Nordic Design Research Conference, 2013*, pp. 406–409., https://doi.org/10.7551/mitpress/11192.001.0001.

Tharp, Bruce M., and Stephanie Tharp. *Discursive Design: Critical, Speculative, and Alternative Things.* MIT Press, 2019.

Transport for London. “Ultra Low Emission Zone.” *Transport for London*, https://tfl.gov.uk/modes/driving/ultra-low-emission-zone. Accessed 8 December 2021.

Weatherall, Jo White Ann. “A Grounded Theory Analysis Of Older Adults And Information Technology.” *Educational Gerontology*, vol. 26, no. 4, 2000, pp. 371–386., https://doi.org/10.1080/036012700407857.

Works Consulted

"Vehicle-to-Vehicle Communication." *NHTSA*, 2017. https://www.nhtsa.gov/technology-innovation/vehicle-vehicle-communication. Accessed 14 July 2021.

Abrams, Janet, and Peter Hall. *Else/Where: Mapping New Cartographies of Networks and Territories.* University of Minnesota Design Institute, 2008.

al Rabbaa, Jad. *MRSive*. OCAD, 2019. http://open-research.ocadu.ca/id/eprint/2496/35/AlRabbaa_Jad_2019_MDES_DF_Thesis.pdf. Accessed 1 June 2021.

Ayoubi, Ayda. *Starling Crossing.* Umbrellium, 2017. https://umbrellium.co.uk/projects/starling-crossing/

Certeau, Michel de. *The Practice of Everyday Life*. University of California Press, 1984.

Dunne, Anthony, and Fiona Raby. *Speculative Everything: Design, Fiction, and Social Dreaming.* Cambridge, Massachusetts: The MIT Press, 2013. Print.

Hillier, Bill. *The Social Logic of Space.* Cambridge, 1984.

Lefebvre, Henri. *Critique of Everyday Life*. Verso, 2008.

Mollerup, Per. *Wayshowing > Wayfinding: Basic and Interactive*. BIS, 2013.

Pritchard, Tom. “Self-driving cars: Here’s what the autonomous driving levels mean.” *Tom's Guide*, 2021. https://www.tomsguide.com/reference/self-driving-cars-heres-what-the-autonomous-driving-levels-mean. Accessed 10 May 2021.

Reinfurt, David. *A New Program for Graphic Design.* Inventory, 2019.

Soja, Edward W. *Postmodern Geographies: the Reassertion of Space in Critical Social Theory.* Verso, 1989.

The Purdue OWL. Purdue U Writing Lab, 2019.

Appendices

Appendix A: Concept Map from Ideation Process

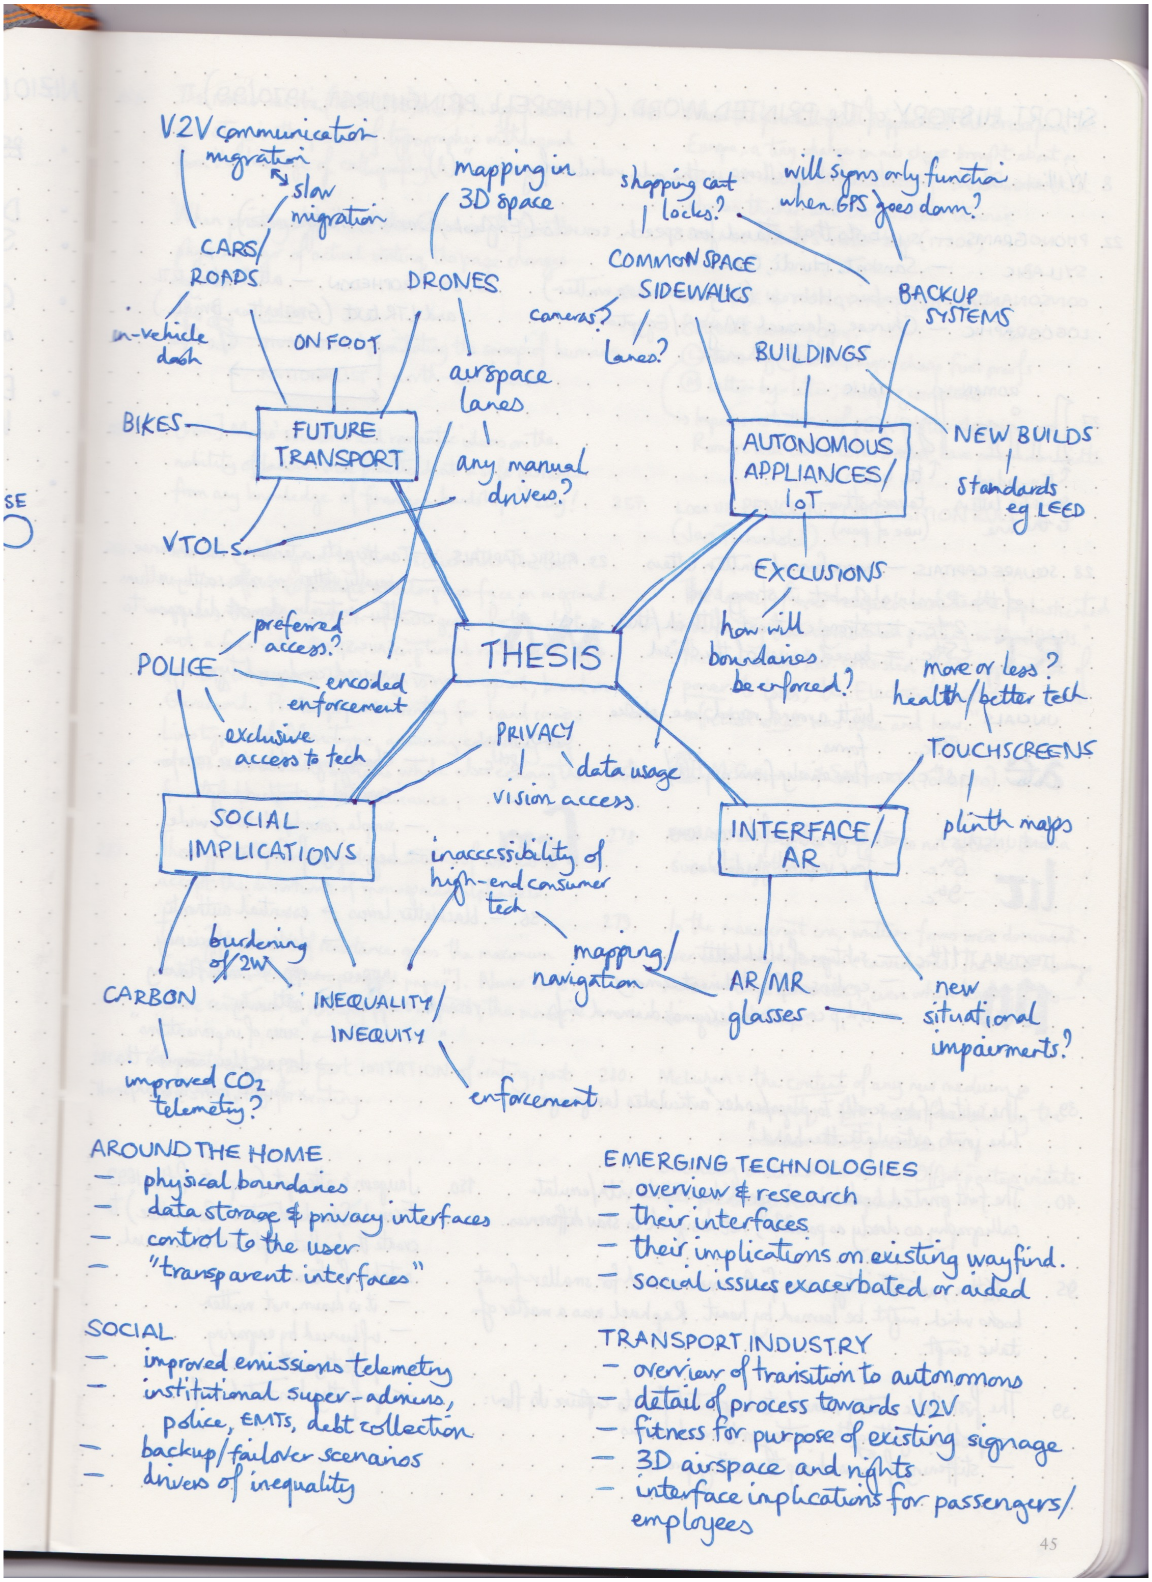

Appendix B: Paper Zines from Research Phase

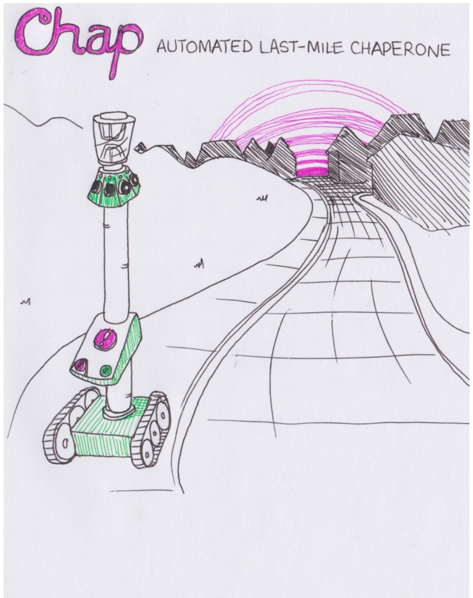

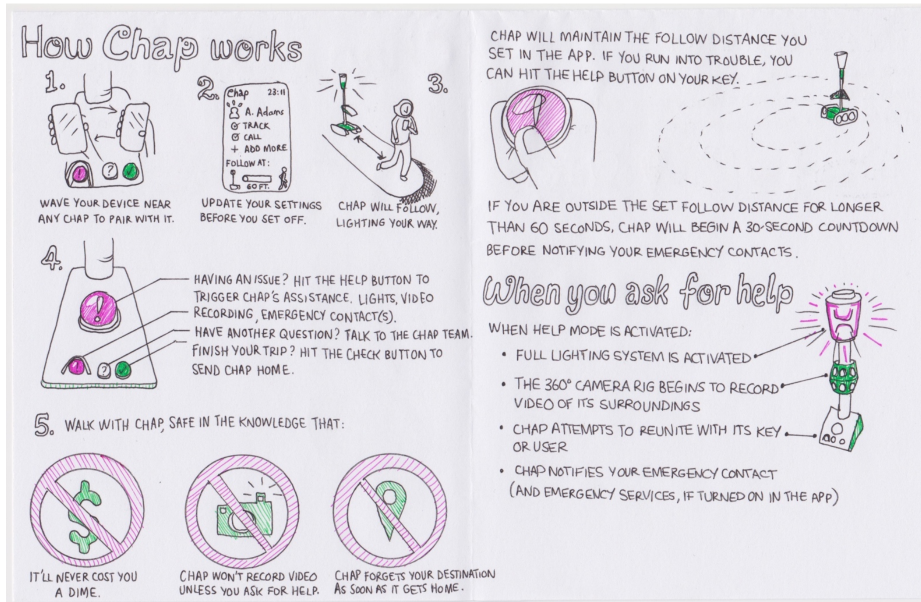

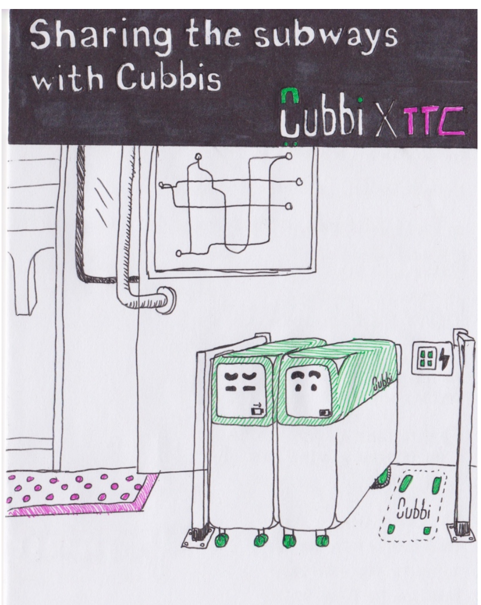

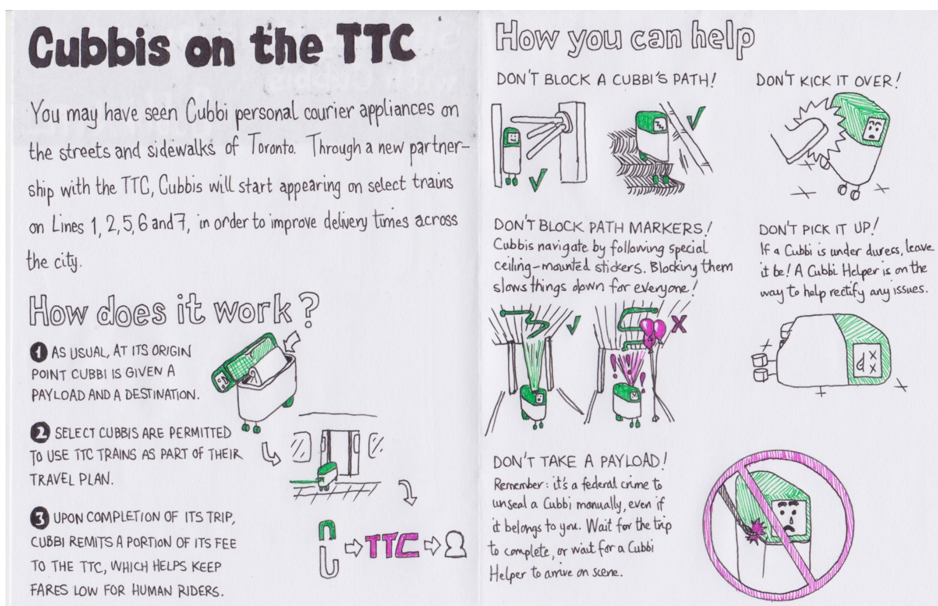

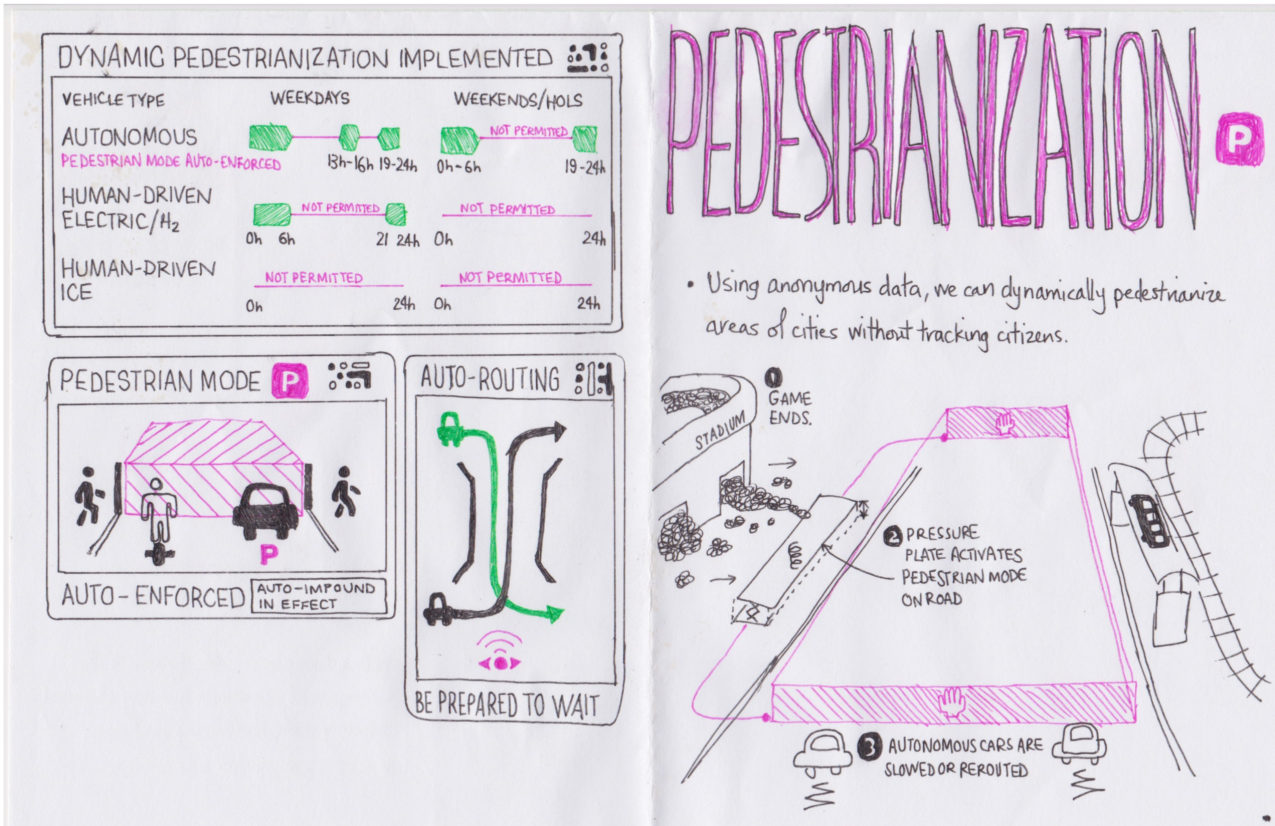
DYNAMIC PEDESTRIANIZATION IMPLEMENTED
VEHICLE TYPE
WEEKDAYS
WEEKENDS/HOLS
AUTONOMOUS
PEDESTRIAN MODE AUTO-ENFORCED
13h-16h 19-24h
0h-6h
NOT PERMITTED
19-24h
HUMAN-DRIVEN ELECTRIC/H_2
NOT PERMITTED
0h 6h
21 24h
0h
NOT PERMITTED
24h
HUMAN-DRIVEN ICE
NOT PERMITTED
0h
24h
0h
NOT PERMITTED
24h
PEDESTRIAN MODE P
P
AUTO-ENFORCED
AUTO-IMPOUND IN EFFECT
AUTO-ROUTING
BE PREPARED TO WAIT
PEDESTRIANIZATION
P
• Using anonymous data, we can dynamically pedestrianize areas of cities without tracking citizens.
1 GAME ENDS.
STADIUM
2 PRESSURE PLATE ACTIVATES PEDESTRIAN MODE ON ROAD
3 AUTONOMOUS CARS ARE SLOWED OR REROUTED

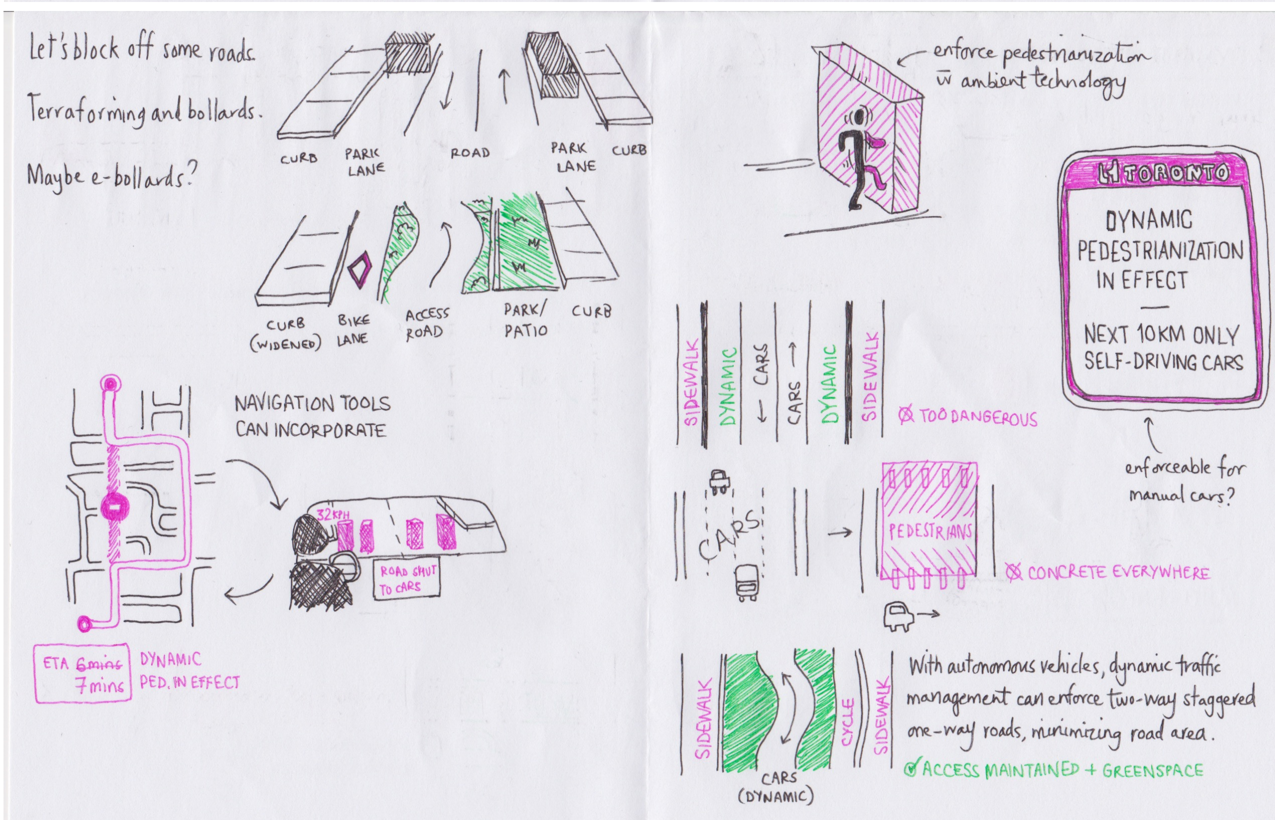
Let's block off some roads.
Terraforming and bollards.
Maybe e-bollards?
CURB
PARK LANE
ROAD
PARK LANE
CURB
CURB (WIDENED)
BIKE LANE
ACCESS ROAD
PARK/PATIO
CURB
NAVIGATION TOOLS CAN INCORPORATE
32KPH
ROAD SHUT TO CARS
ETA 6mins 7mins
DYNAMIC PED. IN EFFECT
enforce pedestrianization w̄ ambient technology
TORONTO
DYNAMIC PEDESTRIANIZATION IN EFFECT
NEXT 10KM ONLY SELF-DRIVING CARS
enforceable for manual cars?
SIDEWALK
DYNAMIC
← CARS
CARS →
DYNAMIC
SIDEWALK
⊘ TOO DANGEROUS
CARS
PEDESTRIANS
⊘ CONCRETE EVERYWHERE
With autonomous vehicles, dynamic traffic management can enforce two-way staggered one-way roads, minimizing road area.
SIDEWALK
CARS (DYNAMIC)
CYCLE
SIDEWALK
✓ ACCESS MAINTAINED + GREENSPACE

Appendix C: Digital Zines as Progress Towards Outcome Refinement

Below: by incorporating time-of-use pedestrianization into city datasets, GPS apps and services can make adjustments during route planning

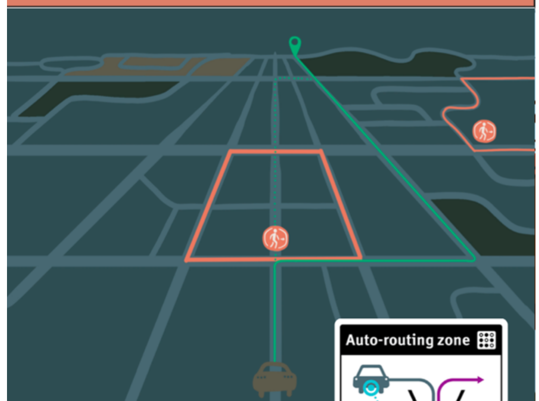

Right: historically wide roadways could provide two-way passage with fewer lanes, by using V2X technology to restrict or impel movement at shared intersections

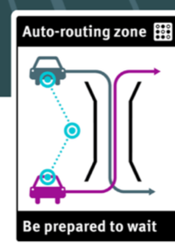

Front: dynamic pedestrianization can fully prevent autonomous vehicles from entering pedestrian space using V2I technology

Above: pressure plates enable prioritization of pedestrian travel at intersections

Right: permitting off-peak use of pedestrian space may incentivize the adoption of vehicles that can process dynamic pedestrianization data

Dynamic pedestrianization in effect

Vehicle type	Weekdays	Weekends and holidays
Autonomous	Permitted overnight, offpeak (12 - 6am; 1 - 4pm; 9pm - 12)	Permitted overnight (12 - 6am; 9pm - 12)
Human-driven electric/hydrogen	Permitted overnight (12 - 6am; 9pm - 12)	Not permitted at any time
Human-driven combustion	Not permitted at any time	Not permitted at any time

Appendix D: Chapter Mockups for Expo 2048 Website

A Changing Toronto

Toronto's cool climate, educated populace and proximity to major finance and industrial hubs ought to give it the edge to thrive during the remainder of the twenty-first century.

Significant risks threaten its growth: a worsening housing and homelessness crisis exacerbated by archaic development laws, and a growing divide between the political aims of the outer suburbs and those of the dense downtown core.

How might the city of Toronto look in fifty years' time? Will work have been done to improve the city's resilience to climate catastrophe? Will automation and the proliferation of technology fundamentally change the urban-suburban relationship? Will densification and infill help to make the city more walkable, navigable, and livable?

Despite rising sea levels, scientists expect increased evaporation from higher ground temperatures to more than offset any rises, likely resulting in a slowly receding lakeshore. The existing Harbourfront may find itself further and further uptown.

Funded and proposed expansions to the vital TTC subway system would give us a much different 2067 system map than today's familiar diagram.

Vital construction projects include the Eglinton Crosstown, the Ontario downtown relief line, and a potential lakefront line to stitch together currently disparate neighbourhoods.

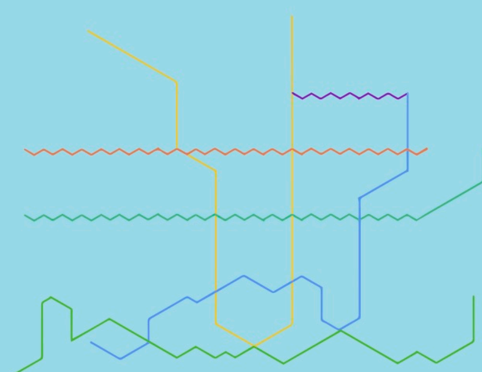

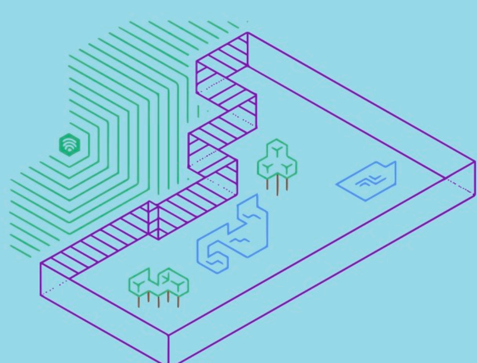

As high-speed data connections become commonplace, expect more and more citizens to seek a thoughtful, if temporary, separation from always-on connection.

Toronto's copious green space and successful greenbelt program may be adapted into an 'off-grid' area where active citizens can truly get away from it all.

Helpful Autonomy

Many twentieth-century positions have been obviated by technological advances. Often, positions of assistance are the first to go: elevator doormen, gas station attendants, toll booth workers. We've lost a lot of helpful jobs. This has implications not only for the future of work, but for the cooperative fabric of our society. More and more, we are expected to fend for ourselves.

History has shown us ill-equipped to adapt to structural or technological unemployment, but as we move beyond industrial automation and into intelligent automation, is there an opportunity to reintroduce a sense of helpfulness? Autonomous machines do not tire, do not require breaks, and are more resilient in the face of poor customer behaviour.

As it arrives, will autonomy operate only as a scythe of disemployment, or can it be used to improve citizens' welfare and safety, reduce wastefulness, and aid in their navigation around cities?

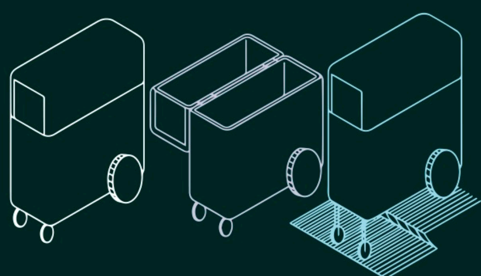

With the evolution and miniaturization of sensing technologies such as LIDAR, computer vision, and object detection, autonomous vehicles will be able to freely navigate through 'free-form' spaces.

The delivery of packages, mail and food will likely occur across a variety of modalities and price points. The cheapest autonomous deliveries will use pre-existing human transport methods like sidewalks, buses and subways. Cubbi uses the TTC as one part of its trip-planning system.

The personal risk of last mile trips - from the bus or subway station to your home - are frequently cited as a factor in people's decisions not to use transit for their commutes.

Autonomy may have a low-cost answer for that: autonomous last-mile chaperones that follow at a safe distance, reacting to the movement of a handheld beacon and remaining ready to call for assistance.

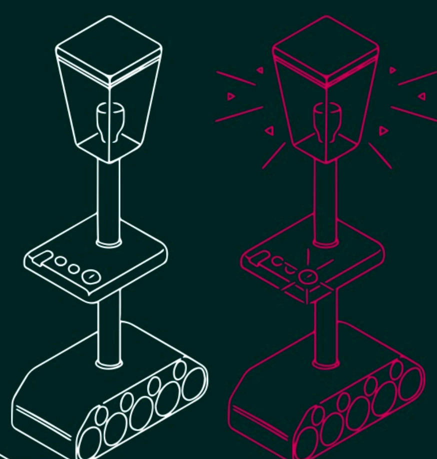

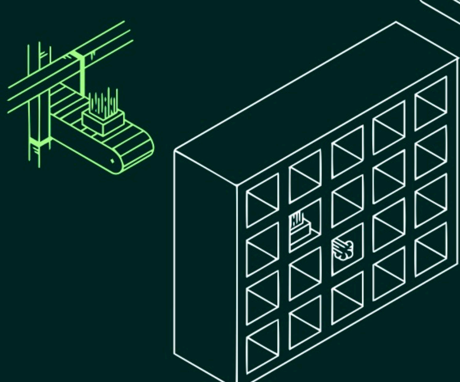

Automats were a popular self-service restaurant in the 1920s, where food and beverages appeared, as if by magic, in individual cubbies, served by a conventional kitchen behind.

A subterranean hydroponic or vertical farming facility could autonomously grow, portion and package small servings of fresh fruit and vegetables in urban environments, reducing food waste, maximizing land use and geothermal properties, and democratizing access to healthy food.

Appendix E: Accompanying Digital Materials

Title, Description, File Name	Date	File Details
Envisioning Autonomy Poster Series *Digital versions of exhibited poster series* EnvisioningAutonomy_Posters.pdf	April 5, 2022	PDF, 12.9MB
Intersection Management – Playthrough Video *Playthrough video of web experience* EnvisioningAutonomy_Intersections_Video.mov	April 13, 2022	QuickTime, 125.2MB
Sharing your Crosswalks – Playthrough Video *Playthrough video of web experience* EnvisioningAutonomy_Crosswalks_Video.mov	April 13, 2022	QuickTime, 97.7MB

Appendix F: Web Experience URLs

NOTE: These URLs may not function indefinitely due to code dependencies and server hosting changes.

Sharing your Crosswalks
https://grgmrtn.github.io/EnvisioningAutonomy/ocadrepo/cross-walks.html

Intersection Management
https://grgmrtn.github.io/EnvisioningAutonomy/ocadrepo/intersec-tions.html