AI ETHICS

AI Localism

By Stefaan Verhulst

SUMMARY KEYWORDS
ai, cities, instance, artificial intelligence, governance, localism, principles, regard, relates, concerns, increased, oversight, citizens, application, order, certain kinds, engagement, context, considerations, decision

00:02
Hi, everyone, my name is Stefaan Verhulst, I'm the co-founder of GovLab, or Governance Laboratory. We are an action research center based here in New York, New York City. And our mission is to transform and improve the way we go about making decisions in the public interest using new technologies and new science. And specifically, what that has meant is that we have focused on two important assets that technology can enable us to use for decision making, on the one hand, people so we do a lot of work around collective intelligence, citizens engagement, and crowdsourcing. And then the other asset that we are focusing on is data. And it is within that context that we do also a lot of work around artificial intelligence. And specifically, one of the areas that we have focused on is the use of artificial intelligence at the local level, and particularly how cities and localities can govern AI in a way that meets international standards with regard to trust and legitimacy. And that's what I would like to focus on with all of you today, which is what we call AI Localism. So let me start sharing my screen to share a few observations towards that end.

01:37
As I mentioned, AI Localism is mainly about the responsible use, and design of artificial intelligence at the local level. And obviously, we have seen a massive increase in the use of artificial intelligence, and especially the experimentation in the use of artificial intelligence worldwide, within a wide range of cities. Whether it's in Atlanta, where Amazon Prime is starting to use robotics and AI for delivery to ultimately in Copenhagen, where we are also seeing the use of artificial intelligence for traffic management. And so what I will try to do first is to highlight a few use cases on the positive uses of artificial intelligence at the local level. And then I will reflect briefly on some of the concerns that automatically means that cities also have to cope with some of those concerns, and as a result, have to start thinking about governance of AI in a particular unique way. And what I hope to convey today is that cities are actually at the forefront of governance. Many nations and national strategies are advocating for AI. But they don't really are embracing governance, nor do they suggest ways to establish trusted users of AI in a more concrete and detailed manner. So we see cities at the forefront of governance. And we see cities becoming innovation hubs in the way we go about regulating, and also governing AI. And to a large extent, that's a result of indeed, as I already mentioned, the increased use of AI within the city context. And we see across the area of city services, but also city concerns, we see that cities are starting to use machine learning and a variety of other artificial intelligence applications to for instance, improve energy consumption, education, safety, social care, housing

Transcribed by https://otter.ai

provision, urban planning, public services, and also engagement with their own citizens and residents. Just going to give you a few examples to give you a flavor of what we are talking about. And so one example is, for instance, in the state of Vermont here in the United States, where Vermont has been using artificial intelligence in order to do enhanced weather analytics, so that for those that are dependent on the weather forecasts, such as for instance, farmers, but also those that might need to think about how do we provide energy and how do we manage the grid for an energy grid for instance, those users are starting to use weather analytics in a far more enhanced way where machine learning is being used to predict weather patterns within the context of Vermont, in order to provide for better for instance, energy provision, and forecasting, as well. We also see the use of AI in order to deal with, for instance, the homelessness crisis. And this is an example, which is called StreetLink, where some of the cities are using AI in order to really manage alerts with regard to homelessness, but also, in order to prevent homelessness of occurring. We also see of course, the increased use of machine learning to improve bus time accuracy, so that you don't have to, for instance, wait in a shelter, that might be dangerous, or might also not be as comfortable as one would hope. And so we see the quality of life, but also the quality of safety going up if we know when the bus will arrive in a more accurate fashion. We also see the use of artificial intelligence, in order to predict fire risk based upon for instance, housing directories. And a certain kind of scoring, that is being done by artificial intelligence, one can actually hone in on those houses or compounds that have a higher risk with regard to fire. And instead of having the fire department visit every home, they can actually really target those places that might have an increased danger for fire. And this is an example of New Orleans. We also see here in New York City, for instance, the use of AI, in order to identify gun violence. And so in a variety of neighborhoods, there are sensors that capture gunshots, and then artificial intelligence is being used to locate where gun violence might take place. And so that allows for a much more faster response to some of the dangers that might be associated with gun violence as well. We also see increased public services being offered to for instance, those that are in need, from a mental health or from a social justice point of view. And so in this context, for instance, machine learning is being used to triage incoming calls, and determine which calls need to be answered faster, and who are the appropriate, for instance, counselors to deal with certain kinds of needs. We also see increasingly that artificial intelligence is being used to engage with citizens in a more targeted and also effective manner. This case, for instance, we have a chatbot, or a virtual assistant that managers request with regard to immigration, for instance, both at the local and the national level. We also see the use of artificial intelligence to actually manage collective input. This is an example where artificial intelligence had been used in order to actually make sense of input received through a public consultation effort. And we also see that, of course, algorithms are being used. And again, quite often, in a positive but also quite often a lot of concerns have already been raised towards that end. But here in New York City, for instance, again, algorithms are being used in order to distribute placement of students in high schools, knowing that of course, there is a scarce amount of spots in certain kinds of high schools, and so how does one go about distributing those in a more fair manner. We also see that artificial intelligence, together with so-called 'digital twins' are being used in order to improve urban planning so that one can actually simulate certain kinds of interventions that subsequently will inform better decision making, and also better public service delivery. And so these are just a few examples on how cities are embracing artificial intelligence, and really, a few examples where artificial intelligence can be a positive force.

10:00

Transcribed by https://otter.ai

Now, obviously, there are also a whole range of concerns associated with artificial intelligence. And those are concerns related with, for instance, increased surveillance. If one starts using, for instance, facial recognition in order to monitor crime, but also discrimination and false positives, for instance, when facial recognition, for instance, fails to actually really identify a particular kind of suspect. We also see increased concern with regard to discrimination. And this is, for instance, an overview of the potential harms from algorithmic decision making that was produced by the former CTO of the United States which all those harms, and those concerns can also, of course, be noticed within a city context. Now, the question, and of course, we also see an increased concern with regard to how, for instance, privacy concerns or privacy violations might happen as a result of the increased use of artificial intelligence.

11:18

And so that leads us to then the second half, and the more and I think, the one that I wanted to really go to of this short lecture, which is how are cities responding to some of those concerns, ie, what are the innovations? And what are the actions that we see at the local level with regard to AI governance. And so what we see, and this is the canvas that we're using of AI Localism, in order to document categorize, and assess some of the actions that cities are undertaking with regard to AI governance. And so we see that it's not one particular kind of interventions. But it's really a package of interventions that cities are experimenting with, in order to make AI more trusted, more legitimate, and also, in some cases, more effective as a result of increased legitimacy, as well. And they range from principles and rights that are embraced by cities, to efforts to increase AI literacy among their officials and their residents, to increase transparency about how AI is actually used so that there is an increased trust in the use of AI that it's not a black box that is making decisions, increased accountability and oversight, increased engagement, both with citizens, but also with key stakeholders within their city. Procurement is also a key tool that is being used by cities in order to steer certain kinds of design features of AI. And then of course, we also see the increased legal interventions, and a set of policies that are specifically targeted for AI at the local level, which all this together is what we call AI Localism, which is the AI governance of the applications of AI. And I've given you a few examples within a city context.

13:35

So what I would like to do now is to give you a few examples and flavors on how that happens in practice. And let's start with some principles that have already been adopted by a few cities. And the first example is, of course, Cities for Digital Rights where a set of cities, ranging from New York City to Amsterdam and others, have embraced a set of principles which they call digital rights, in order to steer the application of digital and technology including artificial intelligence within their city. Also, we have seen a few cities that are taking a leading role in order to develop principles that can be then used by other cities as models of principles and they range from for instance, Barcelona, to ultimately Montreal who also has engaged in principle development or in this context, a declaration on how to use AI responsibly at the city level. Now, what are some of the considerations for cities who want to either join those that have already developed principles or develop a set of principles they themselves? And so the first question that cities need to answer is, of course, do we actually need a new set of principles for AI? Or can we apply existing principles, including existing rights frameworks, such as human rights, for instance, that might be sufficient in order to actually steer the applications of AI? Because the concern quite often is that we are getting a set of principles, indigestion, ie for every new technology, coming up

Transcribed by https://otter.ai

with a new set of principles might by the end of the day, make it harder to actually apply those principles. And more importantly, it becomes kind of meaningless if you have a plethora of principles that never really are embraced or embedded. The other question is, of course, how to combine principles with actually practical cost-benefit assessments? Which means that principles have to be really translated into actually implementation guidelines. And some of those implementation guidelines might be related with concrete assessment of the cost, that an application has vis-a-vis the benefit, because you could have the best principles and you could have the best technology, but if the costs outweigh the benefit, then ultimately that might not be a application one wants to implement within cities. Other question is, of course, it's great to have principles, but how do you actually implement them in a timely fashion? Because what we have seen quite often is that there's a lot of attention to principles, and in the meantime, they adopt applications, and so by the time the principles are endorsed, there are already a few practices that might violate those principles moving forward. So how do we come to principles in an agile way, in order to prevent practices and prevent actually, a certain kinds of path dependency, that might not be easy to circumvent, if one is too late? And then is enforcing and obeying AI principles, the responsibility of the developers, the vendors or the decision makers? So it's great to have a set of principles, but it needs to be associated with actually a real understanding on who falls under those principles, who should obey and enforce them. And so these are a set of considerations with regard to the development of principles within the context of AI Localism, as well.

17:44

The second element in AI Localism, and in the way cities go about governing AI is, of course, an increased engagement with citizens quite often to define the principles, but also to engage them to get a sense on what is the level of comfort, and what are some of their concerns. And so here are two examples. One is around the open dialogue, on AI ethics that took place in several cities worldwide, where cities engaged with residents, and with citizens on what they felt was the appropriate way forward with regard to AI. And then the other example is the one that we at The GovLab have developed, which is called the Data Assembly, which is a citizen's assembly around the use of data and AI in the context of COVID-19. But also, similar kinds of efforts can be conducted, of course, in a more broader fashion. Other examples where we see an increased engagement between cities and universities that reside within their city context. So one example for instance, is from Belgium, where Brussels for instance, is engaging with a set of universities and has created the Kenniscentrum Data & Maatschappij, where the city is actually engaging with experts within the university in order to determine how to best implement and formulate an AI response. And again, the other example is, for instance, the Montreal AI Ethics Institute that also advises and engages with the city and vice versa, around how to go about using AI in an ethical way.

19:44

So what are some of the considerations as relates to engagement that one needs to take into considerationt? Well, the first one is of course, as it relates to segments of citizens, how can we make sure that the engagement is representative, and that the deliberation happens in a way that includes all those that might be affected by the implementation and the application of AI? How do we prevent capture of certain processes by distinctive groups of citizens that might have a particular kind of agenda that doesn't represent the full citizens' consensus with regard to a particular application? As it relates to civil society, what role can civil society and academia play, especially for segments of residents that

Transcribed by https://otter.ai

quite often are hard to engage with? I'm thinking of, for instance, migrants groups or homelessness. So how do we have civil society represent those groups that typically fall out of the and are never invited to a citizens engagement exercise? And then, of course, the other element is not only related to the deliberation, and how we design and structure deliberation, but also, how do we make sure that there is a follow up? How do we make sure that there are some steps that are developed following the deliberation, and that requires that we need to have a really good preparation with regard to impact assessment, stakeholder mapping, and also a good assessment with regard to how we go about designing, for instance, the deliberations. And we would advocate that those deliberations are best designed by actually leveraging scenarios that show citizens and stakeholders on how AI could be used and then get their input on whether that resonates with their values, with their expectations, and whether there are particular kinds of conditions that they would want to see implemented by cities moving forward. So what we have seen is that already two components that are part of the canvas and the package that cities have, in order to define governance. The first one, of course, was related with what are the certain set of principles that should be defined to steer the implementation of AI? And the second one was related with regard to engagement. What we see happening as well, and that's a key element, by the way, related to engagement is that in order to have a meaningful engagement, and to have a city that can embrace, but also can have an understanding on what the potential is, of AI, one needs to also invest in AI literacy. And so we see across the globe, cities investing and creating AI literacy courses for both their citizens and residents, but also for quite often officials. And so here are just a few examples on and especially the one, for instance, that comes out of Finland, elements of AI, a few examples on how AI literacy is being developed.

23:37

Now, clearly, there are a few considerations here as well, which is really about how does one go about demystifying AI without having to go into details that might take a certain level of capacity and a certain level of knowledge that is required to understand actually the cause? And also how does one go about actually touching citizens when it is appropriate to have that kind of literacy. The other tool in the toolbox of cities is of course, laws and policies. And here, we actually do see that many cities are stepping up their legal activity in the context of AI, especially as it relates to for instance, facial recognition, where several cities have already banned facial recognition and which is of course an extreme form of legal intervention and governance. We also have seen that other laws are under construction. For instance, here in New York City, we have seen that there is some legal activity in order to, for instance, audit, automated employment decision tools. So that there would be a requirement that if you are as a city, using AI to screen, for instance, those that are applying for a job, you actually need to have an extra scrutiny, extra audit, in order to understand whether there are biases or whether there is discrimination happening as a result of the use of those automated employment decision tools. And so very interesting that this is becoming the subject of laws, as well. Now, a whole variety, of course, of consideration, both in terms of timing, when does when intervene? Because the big challenge is, of course, it could on the one hand be too late and harm might be done, trust might be undermined. But it could also be too early and you might actually fail to see what is the potential for certain kinds of AI applications in a positive way. Right? And so that's a key concern for any legislation. But it's especially problematic, of course, when some of the technology are still under development, and it might be too early to ban a technology, because it might also involve an opportunity cost as well. Should and if so, how can government's include citizens in the drafting of legislation? I think it's crucial

Transcribed by https://otter.ai

that, of course, similar to previous discussions that we've had with regard to engagement, it is crucial that citizens might have a voice or be given a voice in the drafting of legislation. Other consideration, will government sanctioned AI, and private sector AI be subject to the same regulations? Right, because a lot of the regulations that we see happening within cities are basically regulations that applied to city officials. But what about private sector operating within the city? Should they be at the same standard? Or should there be other standards? Right, a key consideration. And, of course, anyway, related with that is that what's the jurisdiction of the city as relates to regulating AI? Right, in some cases, especially as relates to the private sector, the jurisdiction might not be within the city, it might be national, or it might be the state jurisdiction. And so there is no authority in order to actually even regulate and legislate. And so that's the reason why it's important to think about who actually falls under the legislation, and what's the jurisdiction of the city. And then lastly, of course, do actually legislators have the capacity and the understanding and the sophistication, to know how to go about regulation, which is, again, important to then have AI literacy connected with the legislative process as well. In addition to legislation, in addition to education, AI literacy, in addition to engagement. and in addition to a set of principles, we also see that more and more accountability and oversight mechanisms are being used, including the use of the judiciary. And so we see, for instance, in this context, where the judiciary is actually being called upon in order to make rulings as relates to, for instance, facial recognition. But we also see the creation of a set of working groups or committees in order to oversee or assess the application of artificial intelligence, such as the New York City Automated Decision Systems Task Force, or in Seattle, the Surveillance Advisory Working Group. And so these are examples where new oversight structures in this context, a committee or a working group that can assess and oversee the implementation of certain efforts. And sometimes it works well. And sometimes you see that, like, for instance, in the context of New York City that it actually doesn't work out well. And here in the New York City, for instance, we see that the committee or the task force, then subsequently might actually even submit a shadow report because there there was a feeling that they couldn't implement the terms or conditions because of the lack of support by the city as well. Now, we see also the creation of new positions, like for instance, the Algorithms Management and Policy Officer in New York City and that is an indication that of course, cities are taking the accountability and the oversight more seriously, and that they're creating a point of contact, like an ombudsperson that can be designated to be the point of oversight. And the point that can, anyway also solicit concerns or be the person that receives complaints as well. And so, key considerations with regard to accountability of AI in cities - how is accountability and oversight distributed, ie who actually is tasked to provide oversight? Right? And at what level? Right? Is it a mid-level management, is at top level, is it existing of oversight structures, because what you don't want is to have some competing oversight structures that ultimately might undermine each other. How to connect oversight with a decision process, right. So you could have oversight but then to what extent does that translate into actually new decisions and refreshes policies that are already in place? How to determine remedies and sanctions? Because the key concern, of course, of oversight is that it might be oversight without actually any teeth, or without any enforcement capabilities. And so what are the remedies? What are the sanctions if things go wrong? And that's why judiciary, of course, is a interesting venue, because the judiciary typically has more authorities in terms of sanctions and remedies than typically government oversight bodies or functions. And how will disputes be adjudicated, which is, of course another concern with regard to oversight if one identifies a violation? Well, how will one go about assessing the violation and adjudicating the

Transcribed by https://otter.ai

remedy as such? So these are questions that, frankly, are yet still to quite often be figured out. But these are the questions that cities need to become more sophisticated about as relates to AI Localism.

32:15
In addition to accountability, we also see enhanced transparency. And so this is an example in Helsinki, where the city has developed and is using a website in order to share all the algorithms that are being used within the city. And so you can see an increased transparency of the algorithms being used within the city for the city. Considerations here, what to do with blackbox methodologies, when one actually doesn't know how one came to a particular kind of result, right, that might be misaligned with the intention? So how does one open up the black box? Or how does one make the algorithm and subsequently application explainable? How to establish transparency across the data and AI lifecycle? Because the outcome is obviously the result of what happened prior? Right? And so how, how does one know what is for instance, the quality of the data that was used for, for instance, machine learning? Who made the decision? What is the level of veracity? And so on, and so on. And so we need to establish not just transparency at the analysis stage, we also need to establish transparency across the data lifecycle. And how to retain AI openness when working with the private sector? Because in many contexts, cities are working with vendors in order to develop AI applications. And quite often that means that they have to respect certain kinds of agreements that prevent and protect, for instance, trade secrets. And so how do we make sure that we also can have transparency when vendors are involved, that it's not just transparency with what government does, but also what are the arrangements and agreements for instance, with vendors and also what is the code that is being used? Those are key considerations that one needs to take into account when developing transparency rules.

34:27
And then we have procurement and of course, quite often procurement is forgotten when we think about policy toolbox as relates to governing technology. But procurement is a very powerful tool to steer the design of how technology will be applied and used within a city context. And so this is an example of a ordinance that is steering procurement of technology within Berkeley. So here, of course, the considerations are related with how one goes about actually defining the terms of reference, and is when sophisticated enough to define the terms of reference? And more importantly, is that the right way to experiment with new technologies, because procurement is famously not inducive for experimentation, right? Because one needs to already almost know in advance what one needs. And experimentation, especially as relates to AI, and as it relates to new technologies, is required to even figure out what are the potential applications. And so that's a key consideration, of course, that one needs to take into account. But we'll see the development of more agile procurement processes that might be applicable. And also here are two examples of where one has declarations that steer procurement, as opposed to terms of references that are set in stone as well.

36:14
So we see AI Localism happening, we see that cities are experimenting with a whole variety of governance tools. And the question is, of course, are cities becoming the laboratories of AI governance? What are the pros and what are the ons of AI Localism? On the pro side, the positive side, we really see that AI Localism is actually filling a void. And as I already mentioned, at the beginning, the reason why we see cities stepping up is because quite often, there is no guidance from

Transcribed by https://otter.ai

the national or the state level on how to actually govern AI. And so it's filling a void and cities are stepping in, because they do respond to concerns in a far more agile way, and might be feasible at the national level, or the supranational level. Cities have the level of immediacy, right. So they, they have a much shorter and much more rapid feedback loop on what works, what does not work and how to go about this. And they're also quite often the best in order to engage communities, and to empower communities to have a voice in the process of AI Localism. The con is, of course, is that if every city is starting to develop their own approach to AI governance, we will see massive fragmentation of AI governance, which in itself might undermine trust, and might also actually be negative down the line. We also see that at the local level, we might have even increased capture or certain kinds of interests that might not be the case at the national level. And so that's a key concern that one needs to take into account. And the question is, of course, is there actually capacity at the local level to do meaningful AI governance? Are legislators up to speed on what should be in legislation, for instance? And then, of course, one should not assume that cities are, are the ideal situation, we have also quite often dark municipalism where actually dark forces, which could be corruption forces, are taking place that might undermine the AI governance ecosystem. And so we should actually be aware that, indeed, AI governance could actually go wrong, if we are leaving it up to cities that might not have ultimate accountability and governance structures in place as well. A key element from our point of view is that it has to also go hand-in-hand with thinking about how is it that cities also are developing data collaboratives because to a large extent, the potential of cities to develop AI application depends on the availability of data. And so data collaboratives from our point of view, is the next frontier in order to also steer AI and have the potential to use AI in a responsible way. And that's also an area that we have done a lot of work around.

39:39

With that, let me hold here. And and I hope that has given you a flavor on what we mean by AI Localism, which is really about the application of AI within cities, that can be steered by governance interventions. And I've given you a flavor of some of the governance interventions that we see popping up across the globe at the city level, which we believe is a emerging trend of AI Localism. And I've also given you a set of considerations that one needs to take into account and use for further research, but also for further practice as relates to AI governance at the city level. Thanks for your interest. And I hope this was a extension of what one can think about AI ethics, which is an extension of applying responsible AI within a city context. Thanks so much.

Transcribed by https://otter.ai