

A machine learning approach to open public comments for policymaking

Alex Ingrams

Faculty Governance and Global Affairs, Wijnhaven, Turfmarkt 99, 2511 DP The Hague, The Netherlands
E-mail: a.r.ingrams@fgga.leidenuniv.nl

Abstract. In this paper, the author argues that the conflict between the copious amount of digital data processed by public organisations and the need for policy-relevant insights to aid public participation constitutes a ‘public information paradox’. Machine learning (ML) approaches may offer one solution to this paradox through algorithms that transparently collect and use statistical modelling to provide insights for policymakers. Such an approach is tested in this paper. The test involves applying an unsupervised machine learning approach with latent Dirichlet allocation (LDA) analysis of thousands of public comments submitted to the United States Transport Security Administration (TSA) on a 2013 proposed regulation for the use of new full body imaging scanners in airport security terminals. The analysis results in salient topic clusters that could be used by policymakers to understand large amounts of text such as in an open public comments process. The results are compared with the actual final proposed TSA rule, and the author reflects on new questions raised for transparency by the implementation of ML in open rule-making processes.

Keywords: Public participation, transparency, open regulation, machine learning, algorithms

Key points for practitioners:

- Public agencies can boost transparency in the open regulations process by relying in part on machine learning tools accompanied with descriptions of their comments processing algorithms.
- Machine learning tools could also be incorporated into the open regulations process to save time and costs associated with staff time spent selecting and reading comments. This is especially true of proposed regulations that receive a high volume of comments from the public.
- Machine learning tools cannot be relied upon without managers giving careful consideration to how they will be incorporated into human decision-making processes.

1. Introduction

Big data challenges the information managing capacity of public organisations (Brady, 2019; Klievink et al., 2017) Large amounts of data present a kind of information paradox for democratic governance. On the one hand, the public information perspective has, at least since the writings of English liberal philosopher John Stuart Mill (1806–1873), supported the idea that information is the basis for governments and citizens to engage in the rational information-seeking behaviour that underpins understanding, voting, and public decision-making. On the other hand, scholars such as Peter Dahlgren (2005) argue that in contemporary society we are starting to embark on a path of diminishing returns – even harmful and negative returns – due to the massive amounts of unwieldy, unmanageable, unprocessable, and unverifiable information. Information has a powerful effect on public discussion, and therefore public policy (Bozdag, 2013).

A technological response to big data information burdens are automated forms of information processing relying on automated systems programmed by computer algorithms. Tim O'Reilly (2013) proposed the idea of "algorithmic regulation" to refer to the use of computer algorithms that can automatically process and implement decision-making processes using big data. This concept, which has received further elaboration notably from Karen Yeung (2017), offers a way for organisations to sort through data and make more efficient and effective decisions in allocating public services or processing administrative data.

But, while algorithms offer a technological way of addressing big data, there is a growing body of scholarly literature on the potentially damaging effects of algorithms (e.g., O'Neil, 2016; Lyon, 2003; Yeung, 2018). Public organisations are, like private organisations, increasingly in the firing line for using data in ways that are seen as unethical or undemocratic, and citizens have become more aware of the power latent in algorithms used to process data and the need for more checks to be put on policymakers and technology companies (Chatfield et al., 2015).

Thus, policymakers seek more democratic ways of governing big data. More transparency and accountability of algorithms have been called for by some scholars (e.g., Andrews, 2019; Fink, 2018; Medina, 2015). The message from these demands is that new big data tools need to have the technological sophistication to sustain and maintain large amounts of information, while at the same time delivering the democratic benefits of accessibility, impartiality and legal accountability. However, public administration scholars still lack conceptual models of how transparency and algorithmic tools relate to one another. Without such models, it is difficult for policymakers to practically address the information management while also solving the challenges of algorithmic transparency.

To this end, this paper presents a machine learning model for a transparent policymaking process in open regulatory decision making. Machine learning involves programming a computer application to process data, which in this case are public comments about proposed government regulations. Machine learning applications for public comments are not a replacement for traditional ways of processing public comments, but they can be performed transparently, and can potentially offer a way to make the process even more transparent than a human-only decision-making process. Algorithms in open regulatory decision making is a particularly salient case because, not only do such initiatives involve collecting data from citizens, but the *raison d'être* of the initiatives is to increase openness of government. The question addressed in the research is thus: *how could a machine learning algorithm address the demands of transparency in an open public comments rule-making process?*

2. Theoretical background

2.1. Algorithmic regulation

According to Kitchin (2014) big data applications can lead to black box decision making processes where the critical decision-making principles and rules are hidden from the individuals (technology users, citizens, public administrators) who are impacted by the decision making. But digital technologies can potentially make administrative processes more transparent because they create a digital record, comprised of data, transactions, and the algorithms that regulate these transactions (Jaeger & Bertot, 2010; Lourenço et al., 2017). Similarly, algorithms leave a digital record comprised of information inputs, outputs, and the programming rules that I describe below by explaining decision-making steps of programming that have been put forward in the theory of algorithmic regulation.

The idea of algorithmic regulation was initially set forward by O'Reilly (2013), but only in quite loose terms. A body of research developed by Yeung (2018 and 2017) has made the concept more analytically distinct. According to Yeung (2018, 507), "algorithmic decision making refers to the use of algorithmically generated knowledge systems to execute or inform decisions, which can vary widely in simplicity and sophistication". Yeung does not have in mind here traditional forms of regulation that have statute status and that consist of requiring actors within a particular jurisdiction to conform to certain administrative rules and behaviours in order to avoid punishment. Rather, in a broader sense, algorithmic regulation is about shaping administrative behaviour through the automated processing of large amounts of information designed to achieve a pre-specified goal. Such a goal could be improving humanitarian missions through analysis of meteorological data (Johns & Compton, 2019), prompting law enforcement action on the basis of video data (Young et al., 2019), tackling abuse in financial markets (Treleven & Batrinca, 2017), or, as in this paper, making a public rule-making process more transparent.

Algorithmic regulation is in some ways an inevitable development of organisations now working within the constraints of data rich management and digital technology user environments. Organisations frequently deal with a constant supply of data, which they seek to make sense of by processing it according to pre-set goals and formulas. Behavioural changes resulting – either to the organisation or to its environment – are the product of an algorithmic form of regulation. It shapes behaviour precisely because it is designed according to a set of rules that carry authority. This authority derives from two sources: (1) the rules are designed in such a way that they can be shared and replicated, and (2) they are powerful; that is they have the property of big data in using a high volume of data processed at high speed. For Yeung (2018), closely related to the concept and applications of algorithmic regulation are forms of sophisticated data analysis in real time and with multiple types of unstructured and structured data types that allows for shaping of behaviour according to socio-technical assemblages of social norms and technology practices.

Given the breadth of this definition, many different kinds of algorithmic regulation can follow. Yeung (2018) set forward four major categories that define different approaches to algorithm regulation. According to this schema, algorithmic regulation may be (1) *fixed* or *adaptive*; (2) *real time* or *pre-emptive*; (3) *automated* or *recommender*; and (4) *simple* or *complex*. In the first category, an algorithmic regulation can either be designed to basically process a set of input according to a set of rules that are basically meant to stay the same because the processing task is either very simple or is likely to be a recurring task that will not change, or because the task is designed to require as little human supervision at all. The second category refers to algorithmic regulations that are programmed to either make decisions according to incoming data (real time) or are programmed to only respond in particular ways when particular things happen (pre-emptive). The third category refers to algorithms that either automate final decisions or outputs without the need for human input (automated) or that require a second system (either human or computer) to interpret the recommendation and implement a decision. The fourth category refers to algorithms that function according to a limited number of tasks, classifiers, or processing loops (simple) or others that, to the contrary, operate on many of those levels at once (complex).

Beyond the decision-making rules that are embedded in the algorithms themselves, algorithmic decision-making also has an important machine-to-human component describing how programmers should interpret and use the output of algorithms. For example, Wagner (2016) argued that the decision-making apparatus of algorithmic regulation should be understood on two levels. *First order rules* are embedded into algorithms themselves, and *second order rules* make changes to the outputs/decisions of the algorithms. This distinction is more fundamental than Yeung's, and raises an important distinction between technical development of the specific steps and rules that dictate how algorithms function, and then, secondly, the second order rules that are, in a sense, rules about the rules that ensure that the larger goals (social, political, organisational etc) are being achieved.

2.2. Algorithms and transparency

Algorithmic approaches can also be used to enhance traditional quantitative modelling. Wagner's (2016) distinction suggest that algorithmic methods are not only useful together with traditional methods, but that their value often lies precisely in how their insights are integrated with human decision making. Campbell-Verduyn, Goguen and Porter (2017) describe varieties of algorithms used for public decision making as 'through', 'with', and 'by' forms of algorithmic governance. 'Through' types of algorithmic governance are algorithms that policymakers use to govern. 'By' forms, in contrast, are algorithms that govern for policymakers without the need for human agency. 'With' algorithmic governance is somewhere in between the 'by' and 'through' in terms of the degree of machine/human agency used. Thus, the types of systems in place to apply and use the results of algorithmic regulation are of vital policy importance. In this sense, there are many different ways to use algorithmic regulation.

I would argue that each of these types of algorithmic decision-making pose different sorts of challenges for transparency. Transparency, in general, is the ability for people outside of an organisation to obtain meaningful information about what happens internally (Ingrams, 2016). According Heald's (2006) typology of transparency, there are 'inwards', 'outwards', 'upwards', and 'downwards' forms of transparency representing the organisational direction of the giver and receiver of information that is made transparent. Thus, inwards transparency reveals information from government officials to citizens, and 'outwards' is vice-versa. 'Upwards' transparency allows managers to understand what is going on at the front line of an organisation, while 'downwards' allows the front line to see what managers are doing. 'Through' forms of decision-making are a fixed kind of algorithm where humans decide how to use the outputs on a case-by-case basis. The flexibility of 'through' decision making allows for greater complexity of algorithmic designs, but they could require more supervision and awareness of organisational context as a result. As shown in Fig. 1, 'through' decision making raises transparency concerns related to internal organisational decision-making processes, oversight, and legal and normative goals of algorithms. In contrast, 'by' decision making relies on less supervision but tends to have more direct consequences for end users of services. Transparency concerns for 'by' decision making relate to the individuals who are impacted by algorithms as well as the precise design of algorithmic rules and their outputs. 'With' decision making combines characteristics of both 'through' and 'by' decision making, but the precise form of these characteristics will depend on the degree to which it more resembles 'through' or 'by'.

The type of algorithm tested in this paper is a governance 'through' approach as per the Campbell-Verduyn et al. (2017) distinction, and with attention to the 'outwards' dimension of Heald's transparency. Outwards transparency is the type of transparency that is the main focus of current scholarly and public concern of algorithms. Figure 2 illustrates this approach; the model is defined by a structure of both computer and human processes. Manual, computer-aided, coding is (at least to start with) essential to the development of the model. The main informational concerns needed to satisfy concerns about outward transparency are labelled on the far left of the diagram.

Five steps characterise the function of the algorithm, each requiring different arrangements of computer and human processes: (1) *sourcing identification and accessing*, where the data sources are identified and means of extracting data are built and maintained; (2) *data selection and cleaning* is where parameters for data analysis are set and the data is cleaned to prepare it for analysis; (3) *model fitting*, where specifications are applied to the data and then trained to improve model fit for future applications; (4) *visualisation*, where the results of the analysis are presented in ways that allow meaningful empirical and policy information to be inferred; and (5) *recommending*, where the results of the analysis imply specific decisions.

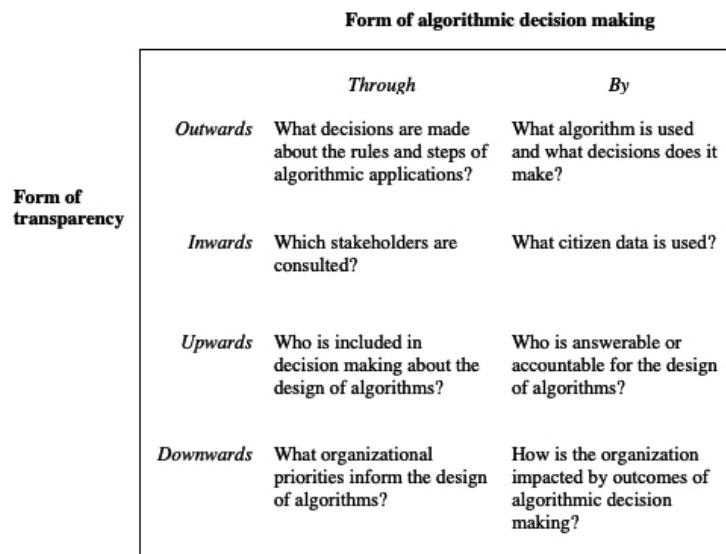


Fig. 1. Transparency challenges in algorithmic decision making.

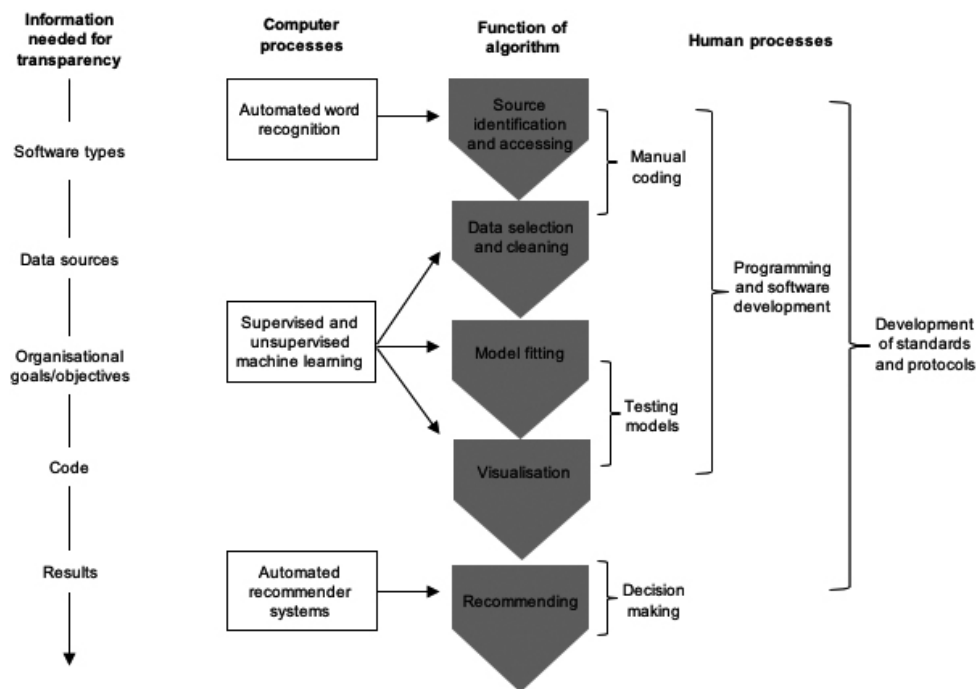


Fig. 2. An algorithmic recommendation model for public comments filtering.

3. Data and methods

In this paper, the author tests such an approach to machine learning using the case of a formal public comments process in the development of new regulations for the use of ‘full body scanners’ in airports

in the United States. In the methodological steps here, the author also demonstrates how a machine learning approach to public comments on regulations can satisfy some of the main transparency concerns for a ‘through’ form of algorithmic decision making with attention to towards transparency. Thus, the software types, data sources, organisational goals/objectives, code, and results of the machine learning approach are detailed in each methodological step. I can be relatively brief on the first three: software types, data sources, and organisational goals/objectives. The latter two (code and results) will need a more extensive explanation. In terms of software types and data sources, the United States federal website, *regulations.gov*, contains full details of new regulations and a docket showing the development – including the public comments process – of the law. *Regulations.gov* also provides developers and researchers with an application programming interface (API) key that can be used to render web content on regulations in formats that can be used for data analysis platforms such as Python and R without the need for web scraping.

The organisational goals/objectives of United States open regulations process are enshrined in the Administrative Procedures Act. There, federal agencies are instructed to seek public input on regulations that are in public interest, to give adequate attention to the comments and explain how they took on board the public feedback in their final rule.¹ The case of full body scanners in airport security terminals is clearly a case in the public interest, and an interesting and important one for many reasons, both empirical and methodological. De Fine Licht (2014) suggests that regulatory topics in spheres that matter greatly to life and death (e.g., health, wellbeing, and financial topics) have not been given adequate attention in transparency research. Empirically, the case of full body scanners in airports is a controversial policy topic that has received a significant amount of public attention. It is therefore one that is likely to attract a substantial amount of public comments from a range of different points of view. The case is also one that has a quantifiable scope in terms of the locations (airports), users (passengers and airport security personnel), costs (numbers and types of technologies), and legal precedents (privacy, data protection, national security). From a methodological point of view, the case is suitable because it includes a large sample of public comments, which is ideal for the kind of text mining used here. Furthermore, the case is also now closed to public comments (so all possible public feedback has now been collected) and the final regulation has already been approved and passed by the United States Congress.

The machine learning analysis is performed on the public comments texts submitted by members of the public, employees of businesses, and employees of governmental organisations (the public comments platform requires submitters to identify themselves according to one of these three categories). The comments were submitted to a notice of proposed rulemaking (PNRM) by the Transport and Security Administration (TSA). Submissions were allowed between March 23, 2013 and June 24, 2013. Submitters could also use a postal address to send comments, though these comments are not transferred to the *regulations.gov* site and the author could not find details of these comments online. Empirical studies of *regulations.gov* have found cases of laws that received thousands of comments automatically generated by bots (Weaver, 2019). The public comments process can also attract organised interest groups who exhort their members to submit a pre-worked text that makes specific policy points. Indeed, in the case of the TSA proposed rule 657 were found to be identical or very similar and were removed during pre-processing of the text corpus. The final number of comments was 4890.

The proposed rule is called ‘Passenger Screening Using Advanced Imaging Technology’ (TSA-2013-0004). The rule essentially adjusts existing regulations regarding what kinds of screening processes

¹The Office for Management and Budget rules on whether a regulation is of interest to the public. But only rarely, for either security reasons or very minor changes to regulations are proposed changes not required to be opened to public feedback.

are allowed at airports. It introduces a gradual program of rolling out advanced imaging technology (AIT) (also known as full body screening) as a primary tool for security screening including details of procurement and deployment procedures. Automatic target recognition (ATR), which is a way for images to resolve to a level necessary for detection while showing only generic outlines of the body, are required in the new regulation. Radiation levels are set within a safe range and privacy safeguards are explained. The latter says that the image will be immediately deleted after the passenger has cleared customs and that security personnel cannot operate personal phones or computers in the screening area. Passengers are free to opt out of the AIT screening and have a 'physical' search instead. The procedures detailed online and on posted signs at airport checkpoints are also set out. Finally, the text of the proposed regulation includes a lot of contextual information about the justifications, costs, and test results that have gone into the policy deliberation process.

The comments were collected using the regulations.gov API. The API gives requesters of web content the ability to select specific parameters of information and collect the information in Hypertext Transfer Protocol (HTTP) to read into statistical software packages. In this case, 'R' was used along with R packages that are designed for text mining. 'Magrittr' was used to read the HTTP code into R. Further cleaning of the texts (the public comments) was carried out using 'tm' and the topic analysis relied on the 'tidytext' package. 'Stopwords' (very common words such as prefixes, conjunctions and 'the'), white spaces, and punctuation was removed. The text was converted to lower case letters to avoid words with upper or lowercase letters being counted as different words. Lemmatisation of the text also grouped similar contractions and tenses of the same word (e.g., scan, scanning, scanned) into a common root word using the 'hash_lemmas' lexicon.

Topic analysis used a form of natural language processing called latent Dirichlet allocation (LDA). LDA analysis is a form of unsupervised machine learning that involves setting text mining parameters for software to analyse groups of text. There are other methods for textual dimensionality reduction, but LDA is one of the most widely used and it has a modelling structure that can be interpreted in an intuitive way. It estimates mixed membership of documents in different topics, and each topic is assumed to be associated with certain words across all documents (Blei et al., 2003; Hornik & Grün, 2011). It thus can be used across a wide range of document types and can capture an important part of the semantic cohesiveness of the entire corpus. Each topic is associated with a list of words that are most frequently associated with those topics in the texts. It also allows calculation of a beta score from 0 to 1 for word-topic density showing the degree to which each topic is a representation of the totality of words in the corpus. A score approaching 1 would suggest that the topics are highly associated with the specific words from that topic.

Three hyperparameters of the LDA model were programmed as model priors. This is done so that the algorithm assumes a density of words or documents in the final topic clusters that is appropriate to the semantic diversity of the texts (Blei et al., 2003). Thus, the number of topics (k) is specified in advance, as is the topic-document density (α), and the word-topic density (β). A lower alpha specifies a higher level of exclusivity of topics per document, while a lower beta specifies higher exclusivity of words per topic. The choice of hyperparameters should be informed by the researcher's substantive knowledge of the case, as well as by estimating the influence of different levels of k on the coherency of the results. In the case of public comments on TSA-2013-0004, numbers of texts (comments submissions) vary considerably between different comment categories. Indeed, diagnostic testing using goodness of fit metrics with the 'ldatuning' package was used to identify k and showed that the best range for k was between 20 and 30 (full results in Appendix 2). I decided to use $k = 20$, which is widely used parameter in analysis of similar sized corpora (e.g., DePaula & Harrison, 2018; Walker et al., 2019; Zuo et al., 2019), and avoids the risk of model overfitting by requiring a larger amount of semantic grouping. For β , I followed the

Table 1
Topics, policy area, and top ten most associated words per topic

Topic	Policy area	Top 10 most associated words*
1 Cost-benefit ratio	Costs	Money, cost, security, AIT, machine, safety, benefit, risk, device, increase
2 Invasive technology qualities	Privacy	Advance, technology, image, passenger, TSA, screen, use, see, invasive, oppose
3 Security-liberty trade off	Security	TSA, security, use, machine, safety, little, theatre, public, good, liberty
4 Health implications	Health	Radiation, scanner, use, health, x-ray, machine, expose, cancer, risk
5 Travel logistics	Service	Body, scanner, search, TSA, pat, go, time, machine, scan, agent
6 Waste of public funds	Costs	Waste, scanner, money, TSA, security, taxpayer, time, dollar, make, tax
7 Need for screening	Service	Technology, search, passenger, TSA, screen, use, AIT, detect, alternative, select
8 Authority of the TSA	Civil liberties	Image, TSA, right, rule, policy, public, travel, citizen, agency, propose
9 Privacy violations	Privacy	Privacy, invasion, security, use, right, violation, AIT, machine, personal, invasive
10 Anti-terrorism strategy	Security	TSA, security, AIT, terrorist, American, government, liberty, attach, citizen, state
11 Legal protections from search	Civil liberties	Amendment, search, 4 th , right, violate, cause, person, unreasonable, people
12 Gender inequalities	Privacy	Body, scanner, patdown, passenger, TSA, screen, conduct, traveller, gender, physical
13 Feeling safe while flying	Service	Scanner, do, get, safe, go, fly, feel, people, make, like
14 Stopping individual terrorists	Security	Passenger, TSA, security, AIT, terrorist, plane, people, attack, airplane, door
15 The effectiveness of the scans	Service	Metal, detector, scanner, screen, use, pat, go, time, down, opt
16 Treatment by TSA agents	Service	TSA, use, people, scan, like, agent, good, one, think, just
17 American public values	Civil liberties	US, country, people, make, will, American, freedom, can, good, one
18 Inconvenience of scanners	Service	TSA, screen, security, time, line, people, airport, travel, agent, wait
19 Sense of privacy and exposure	Privacy	Body, full, scanner, TSA, use, nude, concern, prove, safety
20 Risk of sexual assault	Privacy	TSA, right, fly, child, assault, sexual, American, government, touch, citizen

*Topic words listed in order of beta score.

convention of 0.1. While this is a common standard in LDA analysis, it is also a standard that suits the structure of a corpus that includes a large number of a relatively small comments texts where each text is assumed not to have many words in common with a majority of the other texts. I also followed the recommended approach of setting α at $50/k$ (Griffiths & Steyvers, 2004).

4. Results

4.1. Machine learning model and in-depth reading of comments

Table 1 shows the results of the LDA topic analysis with the top ten most strongly associated words in the topics. LDA does not automatically give its users a thematic label or description of the topics. Rather the researcher must use empirical knowledge of the texts and contextual sources to determine their own labelling system. At face value, all of the topics bear a semantic resemblance to the TSA law, which is a positive sign that the analysis has identified the kinds of topics that we would expect. There also does not appear to be an irrelevant or vexatious material that could suggest that spam or irrelevant commentary was being included in the corpus of commentary texts.

However, interpretation of the themes requires closer inspection of the public comments to make sense of some words that require context to understand. From the most strongly associated comments for each of the topics, I randomly selected 20 of the comments (400 in total) for closer reading. This reading

informed the labelling approach to the topics as well as improved the substantive knowledge of what kind of public concerns and questions were contained in each topic. This is an important research step but also essential for agency decision-makers to do to improve their understanding of the comments when they are addressing the public input for the revised versions of the rules. Below, I comment further on the meaning of the five topics that were the most salient in terms of having the highest number of most highly associated comments. To make it easier to present the results of the in-depth reading of the comments texts, the topics were grouped by the author according to their major policy areas. There are six policy areas: health, cost, privacy, civil liberties, security, and service.

4.1.1. Health

The health policy area includes just one synonymous topic: *health implications*. Comments in this topic are about the dangers of the radiation effects, and the need to control the amount of radiation that they use. Another area of attention is protecting the health of vulnerable travellers such as children, the elderly, and individuals with medical conditions. Comments also underlined the need to make sure that health experts have been properly consulted about the risk to the scanners being carcinogenic. Another, related concern here is that the rule address the fear of health risks even if the actual risks are low.

4.1.2. Costs

The *cost-benefit ratio* topic relates to the relative costs of other types of services that could also save lives, perhaps more efficiently than the full body scanners. Some comments addressed whether there is sufficient evidence that the scanners could have performed better than earlier technology. There are different kinds of cost trade-offs discussed such as the short-versus long-term benefits of the scanners. Some comments suggested that the costs are disproportionate for their intended benefits and asked whether there is enough proof that they work or that the other problems such as privacy risks are worth the costs.

The *waste of public funds* topic concerns the efficient use of tax money and the need to compare with other forms of airport surveillance in order to really understand the benefits of full body scanners. Some comments question whether the scanners should be used at all and ask whether a good anti-terrorism strategy to address the causes of terrorism would be better than a bureaucratic approach with new technology and daily controls on people. According to some comments, the existing pilot does not warrant a full roll out.

4.1.3. Privacy

The topic about the *invasive technology qualities* of the scanners has some semantic overlap with other topics such as *privacy violations* and *sense of exposure*. These are concerned with what measures are in place to prevent breaches of privacy and constitutional rights and the laws that could be brought to bear to prevent violations. The invasive technology qualities policy area address broader concerns about technology ethics, employing terms such as “virtual strip search” and “x-ray vision” and raising the matter of whether the human-technology aspect of the scanners such as how the scanning settings influence the ability of TSA officials to encroach on body privacy of travellers.

The *privacy violations* topic focus more specifically on the nature of the privacy encroaches, and use words such as “invasion”, “intrusion”, and “violation”. Comments express a concern that the scanners cross a certain privacy line, the matter of the trustworthiness of TSA officials, and the issue of whether the existing pilot for the scanners has sufficiently addressed the possibility of using alternatives that are not so much a threat to body privacy. In contrast, the *sense of exposure* topic, gets into the disturbing possibility

of misbehaviour and misuse by TSA officials, and the lack of protections in current implementation of the scanners for the feelings of embarrassment and humiliation among travellers.

Two other topics in the privacy policy area are also quite similar: *gender inequalities* and *risk of sexual assault*. Gender inequalities express concern over protections for women and transgender travellers in particular, and suggest the current rule is too vague about how these problems will be addressed and that the setting of the scanners may allow too much detailed imaging to be made of private body parts. The risk of sexual assault topic also includes comments that relate more to the TSA officials and the possible misuses that could be made of the scanners by “sexual predators” or “perverts”.

4.1.4. Civil liberties

The *Authority of the TSA* topic includes comments about whether the current system is legitimate. There are several different ways that this legitimacy and authority is questioned. There are questions about whether the agency is “against public opinion” in its methods and whether it shows adequate respect and concern for people. In sum, this topic raises the problem of whether the use of the scanners will lead to security systems being abused. The *Legal protections from search* topic gets into the specific laws that may be at risk. It questions whether the rule is constitutional because it restricts freedom and privacy, and potentially encroaches on religious violations. The 4th Amendment in the United States Constitution is raised in particular in so far as personal searches normally require a warrant, and some comments suggested that it was not clear that scanners in airports or the alternative pat downs also merited a higher level of approval.

Relatedly, the American public values topic seems to get at less legal but perhaps more fundamental value questions such as protections against spying and ensuring that the scanners do not feel like an Orwellian approach to security. There is discussion about the American national sense of ambition, and also a perception that the airport security measures of other countries have not yet been properly studied.

4.1.5. Security

This policy area has three topics: *the security-liberty trade-off*, *anti-terrorism strategy*, and *stopping individual terrorists*. The *security-liberty trade off* concerns the matter of whether the scanners go too far in addressing security. Comments in the topic suggest that there will be a negative effect on liberty and that the current rule doesn’t protect long-term liberty of Americans, given that it might undermine the rule of law and that, in reality, the scanners may have flaws that allow terrorists to hide weapons. Chemical weapons are mentioned specifically as an emerging threat that the full body scanners do not seem well designed to detect.

Anti-terrorism strategy, and *stopping individual terrorists* appear quite similar in their word associations but the former is more about the overall American policy towards terrorism suggesting that the whole war on terror is a distraction because there are other dangers that are more important to people’s lives, and that the country has become preoccupied with “security theatre”. In contrast, stopping individual terrorists has more to do with the lack of adequate testing for how the scanners can be circumvented. Other comments suggest that the scanners are necessary but that they may just shift the focus of terrorists to other targets than airports.

4.1.6. Service

This policy area covers six different topics. *Travel logistics* is about matters of delay resulting from the apparent high rate of false positives caused by the scanners, long waiting lines, and delays that are caused by metallic replacement body parts as well as the length of time that body pat downs take if suspicions

are raised by the scanners or if travellers elect for the pat downs because they do not want to go through the scanners.

Better to train observers. The *need for screening* questions the body of evidence on which the program is based. Here, many alternative ideas are put forward and the comments raise questions about whether such alternatives (such as behavioural monitoring, walk through detectors, and pre-screening processes) could work just as well. The topic of *feeling safe while flying* includes comments about feelings of nakedness and discomfort caused by the scanners and worries about the challenges for individuals with disabilities and people in wheelchairs. The effectiveness of the scanners topic raises specific challenges about whether the scanners actually will deliver the service that they are supposed to and adds concerns about service quality problems for certain people such as women (who should have the option to be attended to by female TSA staff) and frequent flyers who will have to endure the scanners more regularly. The *inconvenience of the scanners* topic is similar to the *feeling safe* topic, though it also gets into matters of long lines and potentially hazardous situations created by long waiting lines, and the perceived excess of discretion that TSA officials have to direct travellers to further pat downs. Finally, the *treatment by TSA agents* topic goes into the aspect of individual TSA agents even further. Comments here suggest that TSA efforts should be focused on hiring more skilled employees and training in racial and religious bias.

4.2. The final TSA rule

The paper concludes here by looping back to the paper's research question and revealing what ultimately happened in the final TSA regulation, and how this compares with the most salient topics and words from the analysis. The results of the LDA analysis and the examination of comments reveals key topics that agency decision-makers can take into consideration when drafting final rules. Additionally, results of the analysis can be published on agency websites to reveal how the final rule does or does not pick up the main points that came out the analysis of public comments. Table 2 shows a summary of how the main issues in the proposed law were (or were not) addressed in the final version of the law. If the issue was explicitly addressed in the final law in such a way that a concrete corresponding change was made (e.g., changing legal definitions or altering numbers of scanners or the costs) then a check mark (✓) is shown in the table. However, if the final text showed no corresponding concrete change then a cross mark (×) is shown in the table.² I would suggest that, in practice, such a table could be used by government agencies or third parties such as civil society organisations to show transparently and succinctly whether or not each topic in the public comments was addressed in the final version of the law. Through such a relatively simple and easily readable explanatory tool following the implementation of a machine learning approach, agency decision makers can meet demands for greater transparency in open public comments rule-making processes. While, even with this kind of tool, the public would not know how and why the final rule decision was made, they would at least have a point of comparison with which to judge the final decision and to question it. I also include here a fuller discussion to explain the results in the table and to relate it to the results from the topic and sentiment analysis.

Cost considerations were a major topic area in the public comments. The final law does indeed lower the cost of the program by some \$90 million. However, this change appeared to be a coincidental change in new schemes used to calculate personnel costs. The costs of the program decreased from \$296 million to \$204 million mostly as a result of decreasing personnel costs. New amounts were estimated for years

²A similar version of such a summary table might also be a useful public information tool for the agencies that run the open comments process.

Table 2
Summary of changes in the text of the final law

Issues per policy area	Addressed?
Costs	
Lowers the cost of the program by \$90 million	✓
Cost of “anomalies” ignored	×
Security	
Demonstrates a sound break-even analysis	✓
No fairer distribution of costs	×
Civil liberties	
No attempt to improve civil liberties	×
Services	
Made the definition of AIT clearer	✓
Statutory language on propriety of ATR use	✓
Privacy	
Addition of gender Transportation Security Officers (TSOs)	✓
Health	
Possible health risks detailed and dismissed	✓

2008–2017 (2.2 billion) rather than the 1.4 billion 2012–2015. The new rule did respond to concerns about cost effectiveness. It had a break-even analysis showing how successful in detecting threats the technology would need to be in order to be cost effective in the long time given a predicted number of future terror attacks.

The final law also made some steps towards addressing costs from false positives or “anomalies”. The new text stated that in direct response to the comments about the vagueness of the word “anomalies”. It also made the definition of AIT more explicitly to be “a device used in the screening of passengers that creates a visual image of an individual showing the surface of the skin and revealing other objects on the body”, rather than the simple originally proposed “screening technology”. However, the new law did not recommend a new distribution of costs to reduce the burden to individual taxpayers.

In response to the privacy concerns, the new rule added statutory language about the use of ATR. ATR would be required in all AIT scanners to make sure the bod images were general outlines not details. The allowance of gender Transportation Security Officers (TSOs) was added. Language for the pat-down procedures were clarified to be in a “private screening location”. However, concerns about governance of accountability were not given new attention except to articulate the existing legal rights that citizens already have.

Health risks were discussed at length in the new law, as a direct response to the concerns raised in the public comments. However, the discussion was mostly a refutation of the concerns voiced about risks from radiation. An extensive discussion of scientific research was produced showing that the health risks are extremely small. In sum, while many areas corresponding to the machine learning analysis did receive attention from policymakers in the final law, many areas were only superficially addressed or not addressed at all. This notably affects the governance of accountability, legal and organisational responses to privacy concerns, and fairer distribution of costs.

5. Discussion and conclusions

The research question posed in this paper was how a machine learning algorithm could improve transparency in an open public comments rule-making process. The above analysis could be used to guide interpretation of the texts in a more transparent way. This serves, to use the language of Heald

(2012) a horizontal transparency (for citizens). Not only does a machine learning approach to open public comments allow for a set of programable steps to be developed for processing the comments, but it also could potentially enhance transparency of open rule-making process by addressing the demands for information discussed above: software types, data sources, organisational goals/objectives, code, and results.

Can this approach to regulatory decision-making solve a challenge of the information paradox, while also delivering transparency? We cannot tell from the analysis in this paper alone whether this will deliver positive impacts for the accountability and legitimacy of public deliberation processes per se as that would require surveys of public attitudes. However, the approach explored in this paper certainly shows that it is possible to deliver more transparency in the form of informative and detailed insights into the opinions of citizens, civil society organisations, businesses, and other members of the public who take part in public comments initiatives.

The application of a machine learning algorithm to analysis of public comments in a rulemaking process produces deeper insights into the meaning of the texts. These insights are not intrinsically superior to human reading of the material, but they do offer a new kind of insight that can deliver better understanding for policymakers as well as rendering the translation of public comments into changes to regulations in a transparent way. A corpus of textual information, such as the collection of public comments, can be rendered in visualisations showing latent factors, keywords, and evaluative sentiments. This approach to algorithmic regulation in a public policymaking process is not a replacement for manual processes of reading, coding, and analysing texts, but it can play a complementary role. Additionally, text mining has several advantages over traditional methods. The first advantage is that machine learning algorithms can ‘read’ a far greater quantity of material at a high speed. The example in this paper involves over five thousand texts, but even ordinary desktop applications can process much higher quantities. The second advantage is that the analysis can apply sophisticated statistical analysis to generate new insights that do not appear through surface reading by humans. Using the analysis in this paper, I was able to show evidence that the final law on full body scanners fell short of the public expectations expressed in comments. Whether or not decision makers were justified in their approach is a separate question. The law (embodied in the Administrative Procedures Act) does not compel decision makers to adopt all suggestions in full, but only to demonstrate that they have given the comments proper consideration. Insights from machine learning tools such as these provides an evidence-based way of showing where decision makers were faithful or not and invite discussion about the final decisions that were made.

The third advantage is that machine learning algorithms can be created in the open, providing a transparent way to manage public participation in public policy. I also see five potential strengths of this approach in terms of transparency in particular. First, the pressure to show responsiveness to the public (enshrined in the US context in the APA) compels decision makers to give more attention to demonstrating publicly how final rules evidence responsiveness to the public comments. Second, members of the public can compare the data from the ML analysis with the final rule. If they are satisfied with the result, this may lead to the government gaining legitimacy. On the other hand, if they are unsatisfied, they can protest about the result. Third, the opportunity to compare the ML analysis with the final decision could lead to attention towards shortcomings in the human decision-making process. This might compel public agencies to improve the way they make regulatory decisions. Fourth, by visualising legal concepts and organisational decision-making steps, the output of the analysis serves as an organisational learning tool. Finally, the ML programming code can be copied and reused by other organisations in their attempts to understand the views of important organisational stakeholders.

However, there are also potentially some downsides of this approach for transparency. According to Coglianesse and Lehr (2017), algorithmic regulation in government raises a risk of undermining “principles

of reason-giving and transparency”. I see four main weaknesses. The first is more of a technical limitation, but it could lead to transparency being undermined. This is that, rather than removing the problem of human bias from policymaking, machine learning actually makes the human component of decision-making even more salient and critical. Most machine learning approaches depend to a high degree on manual coding or initial decisions about important parameters. These processes can be automated, but the question remains of how human interventions in algorithmic regulation interact with the automated processes. There is even a risk of important public feedback being handled by automated processes that may have been set up poorly, or, even worse, may somehow encode democratic deficits and unjust political systems. Second, the possibility of a permanent and transparent public record may cause decision makers to adjust their decisions in ways that are not good for the regulatory process. Third, the flexibility of decision makers to use informal channels for expertise gathering may be crowded out if all decisions are expected to be transparent. Fourth, a ‘datafication’ bias in decision making may take place where the algorithmic inputs shape human decision making in ways that crowds out attention to other signals. This applies especially to signals that are hard to capture in data such as responsiveness to broader economic or political problems.

It is very difficult to make any recommendations a priori for addressing these weaknesses. More empirical research is needed. In particular, there are always organisational decision-making and political challenges associated with the implementation of new technologies in government (Gil-Garcia, 2012; Meijer & Thaens, 2018). In this respect, finding a balance that prioritizes the strengths of ML for transparency while minimizing the drawbacks is an area that requires much more attention. Another related area of research for the future, is the challenge of reconciling big data and governmental transparency in ways that improve democratic governance. An assumption in this paper is that big data analysis can be a positive tool for transparency, but other scholars have questioned this (e.g. De Laat, 2018; Matheus & Janssen, 2020). More empirical research can help to provide some answers to the question of how and when transparency through big data in the public sector improves public governance.

These challenges are attenuated by the fact that algorithms need to take into account both human and machine elements of machine learning in order to make public decisions ethical and accountable (Wagner, 2019). This means that considerations of both human and machine elements of recommender systems and responses are needed, but we have very little theoretical or empirical knowledge of what is possible and when and how such alternatives should be implemented. The present paper has introduced one example of how machine and human elements can interlink for transparency-building progresses and expanded the problematic areas in the machine-human interface that need further attention.

Acknowledgments

The author would like to thank the anonymous reviewers of this article for their careful reading and learned suggestions on improving the manuscript. Their many comments and ideas were important and formative for the final version.

References

- Andrews, L. (2019). Public administration, public leadership and the construction of public value in the age of the algorithm and ‘big data’. *Public Administration*, 97(2), 296-310.
- Arun, R., Suresh, V., Veni Madhavan, C.E., & Narasimha Murthy, M.N. (2010). On finding the natural number of topics with latent dirichlet allocation: Some observations. In R. B. Zaki M. J. Yu J. X. (Ed.), *Advances in Knowledge Discovery and Data Mining*. Vol. 6118, pp. 391-402. Springer, Berlin.

- Berliner, D., Bagozzi, B.E., & Palmer-Rubin, B. (2018). What information do citizens want? Evidence from one million information requests in Mexico. *World Development*, 109, 222-235.
- Blei, D.M., Ng, A.Y., & Jordan, M.I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3, 993-1022.
- Boecking, B., Hall, M., & Schneider, J. (2015). Event prediction with learning algorithms – a study of events surrounding the Egyptian revolution of 2011 on the basis of micro blog data. *Policy & Internet*, 7(2), 159-184.
- Bozdog, E. (2013). Bias in algorithmic filtering and personalization. *Ethics and Information Technology*, 15(3), 209-227.
- Brady, H.E. (2019). The challenge of big data and data science. *Annual Review of Political Science*, 22, 297-323.
- Budak, C., Goel, S., & Rao, J.M. (2016). Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly*, 80(S1), 250-271.
- Campbell-Verduyn, M., Goguen, M., & Porter, T. (2017). Big Data and algorithmic governance: the case of financial practices. *New Political Economy*, 22(2), 219-236.
- Cao, J., Xia, T., Li, J., Zhang, Y., & Tang, S. (2009). A density-based method for adaptive lda model selection. *Neurocomputing*, 72(7-9), 1775-1781.
- Chatfield, A.T., Reddick, C.G., & Brajawidagda, U. (2015). Government surveillance disclosures, bilateral trust and Indonesia-Australia cross-border security cooperation: social network analysis of Twitter data. *Government Information Quarterly*, 32(2), 118-128.
- Clark, S.D., Morris, M.A., & Lomax, N. (2018). Estimating the outcome of UK's referendum on EU membership using e-petition data and machine learning algorithms. *Journal of Information Technology & Politics*, 15(4), 344-357.
- Coglianesi, C., & Lehr, D. (2016). Regulating by robot: administrative decision making in the machine-learning era. *Geo. LJ*, 105, 1147.
- Dahlgren, P. (2005). The Internet, public spheres, and political communication: dispersion and deliberation. *Political Communication*, 22(2), 147-162.
- De Laat, P.B. (2018). Algorithmic decision-making based on machine learning from Big Data: can transparency restore accountability? *Philosophy & Technology*, 31(4), 525-541.
- DePaula, N., & Harrison, T. (2018). *The EPA under the Obama and Trump administrations: Using LDA topic modeling to discover themes, issues and policy agendas on Twitter*. The Internet, Policy & Politics Conference, Oxford Internet Institute, Oxford, UK.
- Desmarais, B.A., Harden, J.J., & Boehmke, F.J. (2015). Persistent policy pathways: inferring diffusion networks in the American states. *American Political Science Review*, 109(2), 392-406.
- de Fine Licht, J. (2014). Policy area as a potential moderator of transparency effects: an experiment. *Public Administration Review*, 74(3), 361-371.
- Fink, K. (2018). Opening the government's black boxes: freedom of information and algorithmic accountability. *Information, Communication & Society*, 21(10), 1453-1471.
- Gil-Garcia, J.R. (2012). Towards a smart state? Inter-agency collaboration, information integration, and beyond. *Information Polity*, 17(3-4), 269-280.
- Griffiths T.L., Steyvers M. (2004). Finding scientific topics. *Proceedings of the National Academy of Sciences of the United States of America*, 101, 5228-5235.
- Heald, D. (2012). Why is transparency about public expenditure so elusive? *International Review of Administrative Sciences*, 78(1), 30-49.
- Hornik, K., & Grün, B. (2011). topicmodels: an R package for fitting topic models. *Journal of Statistical Software*, 40(13), 1-30.
- Ingrams, A. (2016). *Transparency*. In A. Farazmand (Ed.), *Global encyclopedia of public administration, public policy, and governance*. Cham: Springer.
- Imai, K., Lo, J., & Olmsted, J. (2016). Fast estimation of ideal points with massive data. *American Political Science Review*, 110(4), 631-656.
- Johns, F., & Compton, C. (2019). Data jurisdictions and rival regimes of algorithmic regulation. *Regulation & Governance*.
- Kitchin, R. (2014). *The data revolution: Big data, open data, data infrastructures and their consequences*. London: Sage.
- Klievink, B., Romijn, B.J., Cunningham, S., & de Bruijn, H. (2017). Big data in the public sector: uncertainties and readiness. *Information Systems Frontiers*, 19(2), 267-283.
- Lune, H., & Berg, B.L. (2016). *Qualitative research methods for the social sciences*. New York: Pearson Higher Ed.
- Lyon, D. (Ed.). (2003). *Surveillance as social sorting: Privacy, risk, and digital discrimination*. Washington DC: Psychology Press.
- Matheus, R., & Janssen, M. (2020). A systematic literature study to unravel transparency enabled by open government data: the window theory. *Public Performance & Management Review*, 43(3), 503-534.
- Medina, E. (2015). Rethinking algorithmic regulation. *Kybernetes*, 44(6-7), 1005-1019.
- Meijer, A., & Thaens, M. (2018). Urban technological innovation: developing and testing a sociotechnical framework for studying smart city projects. *Urban Affairs Review*, 54(2), 363-387.
- O'Reilly, T. (2013). *Open data and algorithmic regulation. Beyond transparency: Open data and the future of civic innovation*, 289-300.

- Rahwan, I. (2018). Society-in-the-loop: programming the algorithmic social contract. *Ethics and Information Technology*, 20(1), 5-14.
- Treleaven, P., & Batrinca, B. (2017). Algorithmic regulation: automating financial compliance monitoring and regulation using AI and blockchain. *Journal of Financial Transformation*, 45, 14-21.
- Walker, R.M., Chandra, Y., Zhang, J., & van Witteloostuijn, A. (2019). Topic modeling the research-practice gap in public administration. *Public Administration Review*, 79(6), 931-937.
- Weaver, R.L. (2019). Rulemaking in an internet era: dealing with bots, trolls, & "form letters". *George Mason Law Rev.*, 27, 553.
- Weimer, D.L., & Vining, A. (2017). *Policy analysis: Concepts and practice*. London: Routledge.
- Yeung, K. (2017). 'Hypernudge': Big Data as a mode of regulation by design. *Information, Communication & Society*, 20(1), 118-136F.
- Yeung, K. (2018). Algorithmic regulation: a critical interrogation. *Regulation & Governance*, 12(4), 505-523.
- Young, M., Katell, M., & Krafft, P.M. (2019). Municipal surveillance regulation and algorithmic accountability. *Big Data & Society*, 6(2), 2053951719868492.
- Zuo, Z., Qian, H., & Zhao, K. (2019). Understanding the field of public affairs through the lens of ranked Ph. D. Programs in the United States. *Policy Studies Journal*, 47, S159-S180.

Appendix 1. Griffiths, Cao Juan and Arun goodness of fit metrics estimating the best value of k

The Griffiths and Steyvers (2004) estimate has a parameter that is a better fit the higher the number. In contrast, Cao-Juan et al. (2009) and Arun et al. (2010) have the low parameter as the best fit. Each algorithm relies on a different kind of mathematical calculation. Griffiths uses Gibbs sampling to estimate posterior distribution of words to topics at different numbers of topics. Arun uses the Kullback-Leibler statistic to measure divergence between topic distribution and document word probabilities. Cao-Juan minimises cosine similarity between topics.

