

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

Abstract: This workgroup considered whether the policy analysis function in government could be replaced by an artificial intelligence policy analyst (AIPA) that responds directly to requests for information and decision support from political and administrative leaders. We describe the current model for policy analysis, identify the design criteria for an AIPA, and consider its limitations should it be adopted. A core limitation is the essential human interaction between a decision maker and an analyst/advisor, which extends the meaning and purpose of policy analysis beyond a simple synthesis or technical analysis view (each of which is nonetheless a complex task in its own right). Rather than propose a wholesale replacement of policy analysts with AIPA, we reframe the question focussing on the use of AI by human policy analysts for augmenting their current work, what we term intelligence-amplified policy analysis (IAPA). We conclude by considering how policy analysts, schools of public affairs, and institutions of government will need to adapt to the changing nature of policy analysis in an era of increasingly capable AI.

Introduction

Policy analysis is a common function in government in which public servants provide support for decision making with the aim of contributing to better decisions than would be made in the absence of such analysis. While decision makers have always relied on this support, the modern concept of the professional, multidisciplinary, policy analyst operating somewhere between the social scientist and the political operative was articulated by political scientist Harold Lasswell¹ and has endured in much the same form since.

This workgroup was formed around the question *can the policy analyst be replaced by artificial intelligence, and what are the implications for public policy and*

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

*governance should this capacity be developed?*² As AI improves, can we imagine the displacement of human policy analysts by machines that understand questions asked by a decision maker, scour relevant databases for applicable knowledge, apply machine algorithms against diverse data to compute optimal strategies and make predictions about the impacts of various policy options, return an instant briefing, and engage in conversation about the issue with the decision maker?³

This note details the issues that emerged during our workgroup discussions. We start with a description of the current model for policy analysis and identify from there the design criteria for an AI replacement as the basis for identifying the specifications for such a system. Assuming that were developed and deployed, we consider how an artificial intelligence policy analyst (AIPA) would operate in the current model, leading to an assessment of its limitations. Rather than propose a wholesale replacement of policy analysts with AIPA, we focus on the use of AI by human policy analysts for augmentation and amplification of their skills—what we term intelligence-amplified policy analysis (IAPA). Thus, we revise our starting premise by considering how AIPA can serve as a supplement to the traditional human policy analysis competency. We conclude by considering how policy analysts, schools of public affairs, and institutions of government will need to adapt to the changing nature of policy analysis in an era of increasingly capable AI.

What is a Good Policy Analyst?

Policy analysis involves a range of activities including: identifying public concerns and determining their condition; assembling evidence, and analyzing potential public policy responses; projecting outcomes and developing strategies for dealing with trade-offs; constructing and evaluating options for addressing the problem; assembling bureaucratic and civil society coalitions to support policy formulation

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

and implementation; communicating recommendations to decision makers; implementing the collective intention expressed in public policy decisions; and evaluating policies to measure effectiveness or value, and inform future responses to comparable policy problems.

During the first quarter century of the profession, logical positivism was dominant⁴ with a focus on the analysis of problems and potential interventions using techniques such as descriptive statistics, inference testing, modeling, operations research, systems analysis, and cost-benefit analysis becoming staples of the profession.⁵ Despite these advances, debates over the real, perceived, and proposed role of the policy analyst have coloured the profession's second quarter century. Stemming from some high profile failures of quantitative policy analysis to solve complex public policy problems, the post-positivist policy analysis movement called for a balancing of softer skills such as participatory design, stakeholder involvement, citizens' input, and qualitative methods alongside technical mastery.⁶ Graduate programs in public affairs continue to develop their students in a range of skills, but policy analysis remains a core competency for budding public servants⁷ – though whether or not this activity will continue to be valued in practice is less clear.⁸

There is a rich literature on what policy analysts should do, but less recent research on how policy analysts actually operate in practice.⁹ One survey of government policy analysts found that—when asked to rank five policy analyst archetypes (*connector, entrepreneur, listener, synthesizer, technician*) in order of how they understood and practiced their profession—the “synthesizer” archetype (defined in part as “consulting various sources to understand how a problem is conceptualized ... develop recommended ways to deal with the problem”) was most strongly identified

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

with.¹⁰ We began our design considerations for an AIPA starting from the premise that much of the work of a policy analyst today involves this synthesis activity of responding to a question from a decision maker by finding information (which, for the practicing policy analyst today, usually starts with a Google search,¹¹ complemented by social media listening¹² and feeds from news aggregators) and synthesizing it in an easily-digestible form of decision support (e.g., a two-page briefing note). And while the “technician” archetype (defined in part as “locating of primary raw data sources in order to undertake statistical policy research”) was ranked lowest in the above-noted survey, we also note the resurgence of policy analytics as a specialized input into evidence-based policy making¹³ and consider how an AIPA would support data-rich policy analytics in addition to information synthesis activity.

Policy analysts also engage in activities as a *connector* (i.e., developing cross-government and stakeholder support for policy solutions), *listener* (i.e., understanding how citizens and stakeholders feel about a specific policy issue), and policy *entrepreneur* (i.e., considering new conceptualizations of public problems, and developing creative and innovative solutions). Professional and experienced policy analysts develop domain expertise and situational awareness of a range of technical details in governing, the unwritten rules of how bureaucratic systems operate, and the political context that drives public policy concerns. Operationally, policy analysts maximize their principal’s reach by engaging with, being responsive to, and channeling the politician’s perspective. They protect their principal’s interests by asserting the politician’s perspective in internal and external fora, spotting potential threats, and addressing challenges before they become barriers. Policy analysts also interface with constituents, the broader public, stakeholders, interest groups,

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

government colleagues, officials in other governments, and the politician's political colleagues, seeking to eliminate blindspots and highlight opportunities. Lastly, in an era of information abundance, policy analysts serve as a filter, guarding against unimportant information and directing the attention of decision makers towards important information.

Parts of the policy analyst's skill set (especially the *synthesizer* and *technician* elements) seem amenable to a machine approach to analysis and decision support. We focus on these policy analysis archetypes in assessing how an AIPA can be designed to replace the current model of the policy analyst. However, as we discussed these different views of what policy analysts do, we were reminded of what we knew already: other elements (i.e., *connector*, *listener* and policy *entrepreneur*)—and notably, the trust that develops between the policy analyst and the decision maker, which is necessary for an advisor's advice to have an impact on the decision maker's thoughts—are intensely human activities that do not seem likely to be overtaken by artificial intelligence in the foreseeable future.¹⁴ This foreshadows our conclusion, that a wholesale replacement of the policy analyst is not likely in the foreseeable technology environment. This does not, however, mean that AIPA that supports and supplements the work of human policy analysts is not achievable nor useful. However, we propose the idea of intelligence-amplified policy analysis (IAPA)¹⁵ as a near-term goal.

Design Criteria for an AIPA

Assuming the first task of the policy analyst is to synthesize information in response to a question posed by a decision maker, can AI be feasibly designed to replace this activity? We also considered the design of a technician-type AIPA, either as a direct machine interlocutor with a decision maker or as an assistant to a human policy

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

analyst. What are the design criteria for an AIPA in performing these roles?

We start by sketching three use cases for such a system, asking how an AIPA¹⁶ might be asked to perform in a range of circumstances.

An Artificial Intelligence Policy Analyst: Use Cases in Three Orders of Government

Federal: 5G wireless and the hard decision of whether to ban Huawei

Our current standard for mobile networks—fourth generation, or 4G—will soon be overtaken by fifth generation, or 5G, wireless telecommunications. One issue has come to dominate the discussion around 5G: the Chinese telecommunications giant Huawei. The United States claims that Huawei is a threat to national security because of the as-yet-unproven claim that Huawei equipment contains a backdoor, and the data that moves through their equipment could be made available to Chinese intelligence services.

The Canadian federal government is currently considering Huawei's possible role as a vendor in the development of Canada's 5G systems. The government's decision is complicated by a number of factors: Canada's current, tenuous relationship with China; Canada's ongoing challenges with the United States on a number of trade issues; the Canadian tradition of internationalism, and of leading by keeping in step with others; domestic economic considerations; and uncertainty as to whether Huawei equipment poses a risk. If Huawei equipment does contain a backdoor, Huawei's central role in 5G could undermine the entire security of wireless telecommunications networks.

It is reasonable to assume that an AIPA could understand the questions being asked ("Hey Siri Humphrey, should Canada ban Huawei products in its developing 5G infrastructure?") and synthesize from a search of available materials a two-page briefing note on the relevant issues. This note is unlikely to be superior to what a

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

human policy analyst could do, although it would likely be completed more quickly. Neither approach should conclude with an absolute recommendation, given the inherent uncertainty; rather, a small number of options should be identified, with the pros and cons of each. As with the human analyst, the decision maker could not be certain what biases may have influenced the content or emphasis in the briefing note.

As for technical policy analytics, we find little scope for an AIPA. As this decision is essentially binary - either ban Huawei equipment or allow it - there is little to nothing that can be experimented with before the decision. And even after the decision, the impacts are not of a big data variety but are rather significant event types (e.g., a major security breach from an exploited backdoor vulnerability; an escalation in Chinese displeasure should Huawei be banned, or American displeasure should it not be banned) that are better assessed through human analysis.

Provincial: Emergency room wait times

“Hey Siri Humphrey, how can we improve emergency wait times in the province?” asks the Health Minister. To answer this policy question, a wide range of data would be needed, including hospital administrative conditions (with variables such as shift structures, room capacity, equipment, support infrastructure, and expertise), community demographic data and public health conditions, and the availability of alternative health service delivery options such as health advice via telemedicine, community practitioners, and specialist physicians.

Automating the policy analysis in this case would require gathering all of this data, likely held in separate organizations, and creating a cost-benefit analysis for the province, citizens, stakeholders, doctors, and hospitals. The considerations to health

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

of the people in the province, practical implications such as supply of skilled employees like doctors and nurses, and political considerations such as costs and labour concerns would all have to be weighed carefully. Research and answers to this policy question have been proposed by the medical community, but the political implications and decisions rest with the government of the day. Could Siri Humphrey find a way forward?

More likely, the tough political decisions related to cost, changing demographics of the population, and the decentralized implementation challenges have prevented greater policy action in reducing wait times. However, better application and use of data, gathered and synthesized by an AIPA could help inform advisors and politicians as they weigh careful options, and create better opportunities for small-scale policy experimentation to reduce wait times, without the political risk.

One approach to the question would focus on short term adaptation to demand for emergency room service, and the need to reallocate healthcare resources while still serving those in need of emergency care and those who don't have a community-based practitioner. Predictive analytics could help predict and plan for surges that may be imperceptible to hospital administrators, though this information only has value if hospital administrators have enough flexibility to respond to surges and lulls. AI could shift demand away from overburdened hospitals towards others with excess capacity by proactively steering patients based on an initial triaging of their condition and location (perhaps through a check-in app, although this would assume the Health Ministry can require patients to use the check-in app before being served in an emergency room).

But the question also implies a concern for longer-term investment and policy choices that might alleviate the root causes of the problem. Data sharing amongst

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

hospitals, community practitioners, and the Ministry of Health could be used by the AIPA as the basis for policy recommendations, including the synthesis of information from other jurisdictions and settings.

Local: Public transportation options

Transit covers a range of subjects that can stretch the capacity of politicians and their staff, from routing and pricing to integrating mobility innovations into existing systems and tackling systemic inequality in transportation access. Substantive command of transit issues is only one part of managing the brief. At its best, transit policy development brings together technical concerns, public engagement, and maximizing not just mobility priorities but commercial, sustainability, equity, and other interests as well.

From an operational perspective, consider an alternative to how local bus transit services are run. Traditionally, transit planners analyze data (collected through surveys, rider counts, and demographic and traffic studies) and create routes and schedules and assign capacity to meet demand. Prices are set to balance fairness and willingness to pay with what it costs to run the system. Schedules and routes are then occasionally revised based on rider feedback, public consultation, and other data like rider counts. This approach is an example of traditional policy analysis. An alternative approach would be to offer on-demand transit services that use ML and algorithmic dynamic routing to respond to riders' requests for transportation—Uber for public transit, if you will. Such systems have already been deployed.¹⁷

For longer term planning purposes, automating policy analysis could give policy makers a more comprehensive view of the substantive landscape by integrating qualitative and quantitative data, historical performance and context, contemporary population and community trends, changes in demand for and gaps in service,

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

expressions of public sentiment, and indicators of future demand. By having machines carry out time-intensive research and synthesis, policy staff could have more time to engage in public consultations, manage the complicated politics of transport, and develop more nuanced policy options for their principals. When combined with the output of other tools and technologies, staff and politicians could see how decisions will play out in real life in areas such as congestion, housing access and affordability, environmental impact, and job creation, letting them make subtle adjustments or strategic shifts to benefit favored groups or the general public interest.

As noted earlier, policy analysts currently use a standard web search as a part of their information gathering activities. Search is a relatively easy function to replicate in an AIPA. For more robust search functions that are not vulnerable to the business model of search-engine provider or gaming and manipulation of results by external actors (including data poisoning),¹⁸ recreating the policy analysis search function using something like the IBM Watson system¹⁹ would require populating an expert knowledge database that mirrors the information available to the human policy analyst.

While knowledge repositories have traditionally been difficult to populate and keep current,²⁰ the Government of Canada's GCPedia system contains a repository of public servant knowledge²¹ that could serve as part of a database for training an automated search function. However, expert-populated databases face concerns over potential bias that machine algorithms can replicate and magnify. For example, IBM's celebrated Watson for Oncology system is sold as able to radically change cancer care, by analyzing massive amounts of diverse data (e.g., clinician notes,

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

medical studies, clinical guidelines). Its treatment recommendations, however, are based on training by only about two dozen American clinicians whose hand-coded information on how patients with specific characteristics should be treated is used as the basis for Watson for Oncology recommendations.²²

Social media listening can help to understand citizen's preferences, experiences, values, and behaviours in response to an actual or proposed policy tool change.²³ Upon sending a signal into the policy environment (as either a proposed or real change), social media can be monitored to assess the reaction and adapt the signal in response, with citizen attitudes gauged and observed over time.²⁴ Policy-relevant examples of this approach are steadily increasing.²⁵ This approach to social media listening in support of policy analysis can be automated, built on the initial parameters that guide the machine approach (e.g., what terms to search for, how to evaluate sentiment, how to synthesize comments).²⁶

Synthesis is a different challenge. While AI-driven narrative writing²⁷ and document synthesis²⁸ is improving rapidly, policy analysts' work involves deciding what information to filter out²⁹ as well as what to summarize and how.³⁰ Those filtering choices reflect the analyst's knowledge of the decision-maker's preferences, their professional judgments, and their biases. If used as training data, this knowledge, judgements, and biases would all become embedded in the AIPA's approach to synthesis, perpetuating and perhaps reinforcing the initial biases.

An AIPA could also conceivably act as a technician-type policy analyst, able to query a wide range of data sources, independently develop a model of system conditions and causal mechanisms, and predict future dynamics under either the *status quo*

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

policy or a specified intervention. Current policy analytic systems scan environmental conditions and apply an algorithmically determined policy to achieve a specified objective.³¹ As the presence of sensors and devices in the social environment grows, capturing more data across a range of system conditions, opportunities for algorithmic approaches to monitoring and steering will also expand.³²

Challenge: What constitutes relevant data and information?

Ironically, when policy analysts use Google to search for relevant information, they are already using AI. Google's RankBrain algorithm uses machine learning to process search results and provide relevant information to users, which contains its own inherent biases largely influenced by the business model underlying the search engine.³³ The prospect of policy analysts relying on an AIPA to collect and sort information raises concerns about what information is of public relevance and what will become encoded in decision-making.

Embracing computational tools for policy work means subjecting human discourse, political valence, and knowledge to the procedural logic that undergirds all computation. As algorithms select what information is relevant to decision makers, they provide a knowledge logic³⁴ – thereby gaining power over the flow of information and the assignment of meaning to it.

Policy analysts and builders of the system will need to think critically about what data, methods and models are for a given system. These choices are not neutral or objective, and will force policy advisors to make trade-offs. In order to make these trade-offs effectively, technical expertise will need to be strong not only for initial development but also for continuing adjustments after the system is deployed. The

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

role of the policy analyst will need to evolve to meet this demand.

These shifting roles will pose challenges to two key tenets of democratic governments, transparency and accountability. How can you be accountable if you cannot explain why certain information was or was not considered for advice? A neural network such as BERT, developed by Google AI Language, is a state-of-the-art model for natural language processing (NLP) tasks like question answering, and might be the most accurate way to complete a given task, but may not be as explainable as a rule-based or decision-tree method. The appropriate trade-off between explainability and accuracy might depend on the nature of the advice being provided. Regardless of what method is chosen policy advisors should know what information is being given relevance over others.

Just as transparency is an important tenet of democracy, open code and methods are also valued in the software community. But if some citizens better understand the workings of the underlying method, greater transparency might enable them to game the system and effectively gain privileged access to government decision-makers, threatening increased inequality. The trade-offs between transparency and control are thus both important and context specific.

In deciding what data sources feed into an automated advisor system, policy analysts must consider the nature and origins of the data, in addition to judgments of relevance. These decisions cannot be fully neutral, but are always subject to power, politics, bias, and blind spots. How does social media distort what people actually care about?³⁵ Who funds the think tank that produced a report? Security issues around data sources may extend to the idea of data poisoning, where adversaries flood the system with misleading data or misinformation. It might be easier to recognize these concerns when analysts' time is freed up from collecting

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

data and available to question and reflect on decision inputs. To improve their overall contextual awareness, policy advisors of the future will need to work in more multidisciplinary teams, including both data scientists and policy leaders.

Possibilities and Limits: Thoughts on a Feasible Current System

It is our judgment that it is feasible to build and use a synthesizer AIPA today, based on the following procedures and design elements:

1. The AIPA is given a strictly defined question to investigate based on current data; e.g., “What methods can be used to decrease homelessness in Edmonton?”
2. The AIPA performs a search on a series of textual data sources, including government reports, civil society advocacy and analysis reports, peer-reviewed academic journals, and miscellaneous materials sourced through Internet search engines.
3. The AIPA searches the data for text that may answer the given question. This is already feasible, using a BERT (Bidirectional Encoder Representations from Transformers)³⁶ model trained on a sentence-pair question answering task, as is the case in the Stanford Question Answering Dataset.³⁷
4. The text snippets containing relevant information are aggregated, together with the source data/articles, and presented to the policy analyst as a list, similar to the output of current legal e-discovery systems.³⁸
5. The AIPA reviews this information, discards statements it judges irrelevant, and augments these statements with its own independent research. The scoring of each statement as useful or irrelevant is retained as feedback to the system to inform future tasks.
6. A GPT-2-style (Generative Pre-Trained Transformer version 2) generative text system³⁹ trained on articles and past briefing notes is used to generate multiple draft

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

briefing notes based on the discovered statements. It is important to highlight that this text synthesis is guided by the textual structure of briefing notes in the training data, not a human-level understanding of the question to be answered.

7. The human policy analyst selects one draft from amongst several as a starting point for the briefing note. Their decision patterns are recorded to inform further development of the system.

8. A human policy analyst then edits this note, adjusts it for factual accuracy, and makes subtle adjustments based on their emotional intelligence and understanding of the minister. The initial, GPT-2-generated draft and the final submitted note are both saved, to adjust the generative model using transfer learning so future notes more closely approach the style required by the Minister.

9. The Minister receives the briefing note, and provides feedback to further inform the process (as Ministers routinely do with the work of human policy analysts).

Assuming such a system is developed and deployed, how might an AIPA operate in our current political and administrative settings, and what might be its limitations?

These functions of search, synthesis, understanding, and document generation will become more advanced and reliable over time, due to both general technological advances and progressive specialization to the task at hand. As a result, the process will become increasingly streamlined, particularly for simple queries, so human analysts can focus on problems that require more subtle and complex synthesis of information - such as the question whether Canada should ban Huawei from 5G infrastructure. For the human analyst, this process will increasingly prioritize skills related to managing interpersonal relationships, strategic thinking, and understanding their audience.

Although the synthesizer is a dominant archetype in current views of policy analysis,

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

creating briefing notes is a small part of what policy analysts do. Human policy advisors require skills that would be difficult to replicate with an AIPA, such as empathy and emotional intelligence. Human policy analysts must also be attuned to how the decision-maker thinks, in order to appreciate their personal perspective and anticipate their political concerns. They must balance the fragile dynamics of discretion and bias, seeking the equilibrium where the public good meets the public mood. This balancing act is an essentially human process in decision support, because even decisions that clearly appear to be in the public good sometimes require persuading citizens that this is so. Good policy advisors don't just provide the best options to decision makers, but also the frame in which they should be communicated. The best policy analysts also maximize their reach by engaging, being responsive to, and channeling the politician's perspective to constituents, interest groups, government officials, and political colleagues. They assert their perspective in internal and external fora, spotting and tackling threats and challenges, and projecting the vision of their principal.

There are additional limits to what an AI approach to policy advising and policy making can do. Making use of better information requires that ability and willingness to respond to it. Even with better information, sometimes governments lack the resources or flexibility to respond dynamically – like that hospital administrator receiving predictions of ER demand surges discussed above. Experimentation and design research is key to improving delivery, but not everything can be experimentally manipulated. A/B testing on websites is permissible, but performing randomized control trials using social welfare payments is much more problematic,⁴⁰ even if the condition of social equipoise is met.⁴¹

The Intelligence-Amplified Policy Analyst (IAPA): Artificial Intelligence as

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

Supplement (Rather Than Replacement) for the Policy Analyst

Given the limitations and challenges inherent in moving towards an IAPA, rather than propose the wholesale replacement of policy analysts with artificial intelligence, we argue that the foreseeable future of policy analysis will be one where human analysts use AI to augment and amplify their skills—an approach we term intelligence-amplified policy analysis (IAPA). We thus revise our starting premise, considering how AI can serve as a supplement to traditional policy analysis competencies. As such, we foresee an evolution, rather than a complete disruption of the strategic policy function in governments (barring near-term development of artificial general intelligence—AGI—and avoiding speculation about longer-term radical advances in machine intelligence). This evolution will require adaptation on the part of human policy analysts as there will be increased expectations on their abilities. With supplemental and assistive tools at their disposal, policy analysts may be expected to do more with fewer of their colleagues than currently.

Governments, however, should be cautioned that the IAPA is not a short-cut to cost savings: investments in systems and human capital will be required. For the education, training, and career development opportunities for public servants to provide a foundation and framework upon which to build a public service for the digital future, governments, universities, civil society, and public servants (both current and future employees) must work together to build new competencies, skills, and digital literacies (we conclude on this point, below).

Humans have been increasingly relying on external tools such as AI to achieve higher standards of productivity, and the benefits of this collaboration are being felt across a range of sectors. Yet being affected by AI does not mean that human beings will be replaced by technology, but that technology will be increasingly used to assist people in their work. Therefore the need to upgrade human skills align with the

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

requirements needed to work with AI will become essential. Howard Rheingold⁴² argues that human minds being replaced by a machine is not a phenomenon that the world will experience in the foreseeable future. Instead, he argues that the world will very soon experience the ubiquity of intelligence amplifiers, toolkits, and interactive electronic communities (today's social media) and that these tools will change how people think, learn, and communicate. Personal computers were an early manifestation of these amplifiers, opening up new avenues for people to manage the complexities of the modern world. Scaling up to broader usage, an economy can be said to be intelligently amplified when people are effectively trained in using ICTs to enhance their human intelligence. As the use of AI to augment human capabilities grows to sectors such as public administration, it is estimated that by the mid 2030s a third of public sector jobs will be affected by AI with senior officials and managers predicted to be most affected by automation.⁴³

Intelligence amplification (IA) involves the adoption of artificial intelligence (AI) by knowledge workers as a complement to, rather than replacement of, their activities.⁴⁴ Under an IA framework, human cognitive abilities are made more effective by adopting AI as a tool to support and amplify human intelligence. This view emphasizes the impact of advanced technology in expanding the capability of humans to understand complex problem situations and to derive robust solutions to problems more quickly.⁴⁵ With the help of robotics and AI technologies, humans can be relieved of mundane work, leaving employees with more time to focus on areas that require intrinsically human skills such as creativity, ingenuity, and emotional intelligence.

The rapid growth in AI technology is gradually transforming into an essential component of work life today. Public sector organizations are incorporating AI

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

capabilities to deliver services and increase effectiveness. The Government of Canada has already started using AI tools to accelerate human capabilities; its digital operations strategic plan envisions the need for government to manage digital tools and technology to augment human skills.⁴⁶ A recent example can be found in the adoption of AI by Immigration, Refugees and Citizenship Canada to help with the department's backlog of immigration and refugee claims.⁴⁷ The aim is for AI to eventually deal with refugee applications on its own and, if the project is successful, that front-line immigration officials can ultimately use this technology "to aid in their assessment of the merits of an application before decisions are finalized."⁴⁸

Investing in the Future of AIPA and IAPA

We conclude by considering how policy analysts, schools of public affairs, and institutions of government will need to adapt to the changing nature of policy analysis in an era of increasingly capable AI. New skills and training for policy analysts will be required for the effective and knowledgeable use of AI in support of policy analysis. Skills in research design and investigating the data sources that inform decisions will increasingly be required. Additional skills include human-intensive work such as interviewing, engaging, and consulting to generate qualitative data sources. Human-intensive legwork, such as interviewing sources and engaging communities to gather qualitative data, will remain the domain of human policy analysts. However, the analysis of this data will be increasingly assisted by AI.

Policy analysts will be expected to merge traditional research skills with capacity in data analytics. In the near term, this will involve the crafting of new communication languages between specialist in data analytics and subject-matter expert policy analysts—a process known as paired analytics. As data analytics tools become more user friendly, generalist policy analysts will not only need to learn how to manipulate

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

these tools but, more importantly, to know when they are using the tools incorrectly. Anyone who has ever used a statistical analysis computer package knows how easy it is to use a few simple commands to output reams of results that may or may not be meaningful.

Improvements in technical literacy skills among analysts will be required. Every analyst working with AI assistance must be able to understand the caveats and conditions that come with recommendations. Present-day AI education, given to current analysts and policy-makers and added to hiring criteria, will prepare the next generation of analysts for correctly interpreting AIPA processes and outputs. An important competency for policy analysts to develop is the ability to skeptically interpret AI recommendations and communicate their meaning to others. Human policy analysts will also be required to provide a source of accountability and ensure that recommendations minimize bias (subject to the acknowledgement that human policy analysts embody their own biases).

Perhaps surprisingly, the IAPA will have to improve their human interface skills as much as they will need to adapt their computer interface skills. As AI frees up their time, they will need to engage in forms of data collection that is more community and human centered. To illuminate those who may be digitally invisible⁴⁹ to machine approaches, policy analysts will increasingly need to improve their street-level bureaucrat skills.

Governments will need to invest in and improve the data available to AIPAs. Data interoperability and sharing between departments and governments is a major limit to the development of AIPA. This might include making explicit the implicit knowledge that analysts use, or developing methods of quantifying “soft” data not traditionally available digitally and will allow AI to act in more culturally relevant

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

ways.

Finally, successful movement towards IAPA will require acceptance of new processes by institutional and governmental leaders. While they will likely be removed from the day-to-day requirements of knowing how to correctly interpret the output of an AIPA, they must be willing to adapt to new processes for making and acting upon reports.

Next Steps

This note represents our collective response to the question posed at the outset—can the policy analyst be replaced by artificial intelligence, and what are the implications for public policy and governance should this capacity be developed?—largely written during our deliberations at the 2019 Summer Institute on AI and Society. Because of the fast-prototyping nature of the Summer Institute, this note does not perfectly articulate everything that was voiced during the working group’s time together nor what could be accomplished with more time. As should be clear from the foregoing collection of many bullet points and threads of ideas, we have not yet developed a fully-coherent narrative to describe the many intersecting issues that our seemingly straightforward initial premise entailed.

Following from the Summer Institute and this note, we look forward to subsequent rounds of debate and writing that will build on the above, with the following specific outputs:

The working group will produce a 1000 word think piece, for an outlet such as <https://policyoptions.irpp.org>, that speaks to the public policy practitioner community about any developing fantasies that AI is going to displace the policy analysis function, but also how the profession needs to adapt to changing possibilities. We are targeting a first draft by September 30, drafting collaboratively

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

and video-meeting as necessary. A working title is “Relax, policy analysts—AI is not going to steal your jobs. But it is going to change how you work”.

An academic paper for a proposed Special Issue of the journal *Canadian Public Administration*.

A creative piece (written as a script), tentatively titled “Yes, Siri Humphrey”, will imagine a fully autonomous AIPA substituting for the character Sir Humphrey Appleby of *Yes, Minister* and *Yes, Prime Minister* fame.

An academic grant proposal will be developed assembling a larger team to investigate the technical and applied aspects of continued developed of AIPA.

1. Harold Lasswell. “The Policy Orientation” in D Lerner and H D Lasswell, eds, *The Policy Sciences* (Stanford, CA: Stanford University Press, 1951).
2. A subsequent, more speculative, question this workgroup briefly considered was the role of political decision makers in a future where AI decision support yields optimal policy solutions—that is, if we can replace policy analysts with AI, why not replace politicians too? Why would we want or allow political intervention in decision making that is a deviation from an optimal policy solution? (By analogy: if AI can pilot a fully autonomous vehicle, why would we give a human driver any control over steering, acceleration, or braking? And while such a system might still need a passenger to identify a destination, would representative politics be the best way to articulate collective objectives?). While some aspects of decision maker discretion are increasingly being replaced by automation (e.g., Bryce Goodman and Seth Flaxman, “European Union regulations on algorithmic decision-making and a ‘right to explanation’” (2017) 38:3 *AI Magazine* 50, online: <https://www.aaai.org/ojs/index.php/aimagazine/article/view/2741>), we do not foresee the near-term development of an AI system able to fully replace the type of decision making involved in difficult policy choices, and therefore did not pursue

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

this question further.

3. We return to the design criteria of such a system below, but note here the range of AI paradigms that are at play in such a description. To “understand” a question posed by a decision maker (presumably by speaking to a voice-activated device or sending a short text message) and “engage in conversation” with that decision maker implies abilities in natural language processing (NLP). To “apply machine algorithms against diverse data to compute optimal strategies” implies some form of statistical AI, likely machine learning (ML). However, to “scour relevant databases for applicable knowledge” and “return an instant briefing” implies the type of decision-support / expert systems developed under the “symbolic AI” paradigm dominant in the 1980s (see Donald A. Waterman. *A Guide to Expert Systems*. Reading, Mass.: Addison-Wesley, 1986). Research and commercial success in decision-support systems plateaued in the late 1980s, largely due to a failure to live up to expectations and the difficulties in encoding tacit knowledge in computer systems (Daniel Crevier. *AI : the tumultuous history of the search for artificial intelligence*. New York, NY: Basic Books, 1993), though a reconciliation between the two paradigms is possible (Marta Garnelo and Murray Shanahan. “Reconciling deep learning with symbolic artificial intelligence: representing objects and relations.” *Current Opinion in Behavioral Sciences* 29 (2019): 17-23.).
4. Göktu Morçöl, “Positivist beliefs among policy professionals: An empirical investigation” (2001) 34:3-4 *Policy Sci* 381, online: <https://link.springer.com/article/10.1023/A:1012749120909>
5. Beryl A Radin, *Beyond Machiavelli: Policy Analysis Comes of Age* (Washington, DC: Georgetown University Press, 2000).
6. Frank Fischer, *Reframing Public Policy: Discursive Politics and Deliberative Practices* (New York: Oxford University Press, 2003).
7. Pamela T. Dunning. “Why public administration is needed now more than ever:

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

- Advancing the scholarship of teaching and learning in public administration.” Public Lecture delivered at the Public Administration Conference, University of Northumbria 12 September 2018). *Teaching Public Administration* (2019): 0144739418823824.
8. Nancy Shulock. “The paradox of policy analysis: If it is not used, why do we produce so much of it?.” *Journal of Policy Analysis and Management: The Journal of the Association for Public Policy Analysis and Management* 18, no. 2 (1999): 226-244.
 9. Michael Howlett, Adam Wellstead, and Jonathan Craft. *Policy work in Canada: Professional practices and analytical capacities*. University of Toronto Press, 2017. An earlier inquiry in this vein is Arnold J. Meltsner. *Policy analysts in the bureaucracy*. University of California Press, 1976.
 10. Justin Longo. *Towards policy analysis 2.0*. (PhD Dissertation, University of Victoria, 2013) at 67.
 11. Anne White. *Evidence that Works: Building the Canadian Evidence Infrastructure for Social Policy*. Mowat Centre / Munk School of Global Affairs and Public Policy, research report # 176 (University of Toronto, 2018). The question of bias in the search results that are returned in response to the policy analyst’s web search is returned to below.
 12. Yannis Charalabidis, Euripidis N. Loukis, Aggeliki Androutsopoulou, Vangelis Karkaletsis, and Anna Triantafillou. “Passive crowdsourcing in government using social media.” *Transforming Government: People, Process and Policy* 8, no. 2 (2014): 283-308. Panos Panagiotopoulos, Frances Bowen, and Phillip Brooker. “The value of social media data: Integrating crowd capabilities in evidence-based policy.” *Government Information Quarterly* 34, no. 4 (2017): 601-612.
 13. Giada De Marchi, Giulia Lucertini and Alexis Tsoukiàs. “From evidence-based policy making to policy analytics.” *Annals of Operations Research* 236, no. 1 (2016): 15-38.
 14. Keng Siau and Weiyu Wang. “Building trust in artificial intelligence, machine

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

learning, and robotics." *Cutter Business Technology Journal* 31, no. 2 (2018): 47-53.

15. Vannevar Bush. "As we may think." *The Atlantic Monthly*, 176(1), 101-108, 1946.

Howard Rheingold. *Tools for Thought: The History and Future of Mind-expanding Technology*. MIT Press, 1985.

16. In these cases, we personify the AIPA as "Siri Humphrey" (a spontaneous suggestion made at the Summer Institute by one of its convenors, Ted Parson), a combination of the Apple virtual assistant Siri and the character Sir Humphrey, the senior public servant advising the eponymous politician in the television series *Yes, Minister* (Antony Jay and Jonathan Lynn. *Yes, Minister: The Diaries of a Cabinet Minister* by the Rt. Hon. James Hacker MP. British Broadcasting Corporation, 1981) and *Yes, Prime Minister* (Antony Jay and Jonathan Lynn. *Yes, Prime Minister: The Diaries of the Rt. Hon. James Hacker*. British Broadcasting Corporation, 1989).

17. See

<http://www.theverge.com/2016/9/1/12735666/uber-altamonte-springs-fl-public-transportation-taxi-system>

18. See, e.g.,

<https://www.theverge.com/tldr/2018/7/20/17595584/google-search-results-idiot-donald-trump-images-reddit>

19. This concept was investigated recently by the Commonwealth Association for Public Administration and Management in its SmartGov Discovery project, "a digital platform powered by cognitive computing, also known as machine learning. The system is neither a searchable database nor a web browser, but rather a system that continuously learns from searches, curation and interaction with users. It bridges the gap between data quantity and data insights, builds knowledge, understands natural language and provides confidence-weighted responses." See

<https://www.capam.org/offerings/incubator.html>.

20. Jonathan Grudin. *Why CSCW applications fail: problems in the design and*

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

- evaluation of organizational interfaces. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, 85-93, 1988.
21. Amanda Clarke. *Opening the Government of Canada: The Federal Bureaucracy in the Digital Age*. UBC Press, 2019.
 22. Casey Ross and Ike Swetlitz. "IBM pitched its Watson supercomputer as a revolution in cancer care. It's nowhere close." *Stat Magazine*, Sept. 5 2017. <https://www.statnews.com/2017/09/05/watson-ibm-cancer/>
 23. Grubmüller, Verena, Katharina Götsch, and Bernhard Krieger. 2013. "Social Media Analytics for Future Oriented Policy Making." *European Journal of Futures Research* 1 (1). Springer Berlin Heidelberg: 20.
 24. Paris, Cecile, and Stephen Wan. 2011. "Listening to the Community: Social Media Monitoring Tasks for Improving Government Services." In *CHI '11 Extended Abstracts on Human Factors in Computing Systems*, 2095-2100. CHI EA '11. New York, NY, USA: ACM.
 25. Atkinson, Gail M., and David J. Wald. 2007. "'Did You Feel It?' Intensity Data: A Surprisingly Good Measure of Earthquake Ground Motion." *Seismological Research Letters* 78 (3). GeoScienceWorld: 362-68.; Lazer, David, Alex Sandy Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, et al. 2009. "Life in the Network: The Coming Age of Computational Social Science." *Science* 323 (5915). NIH Public Access: 721.
 26. Khaled Ahmed, Neamat El Tazi, and Ahmad Hany Hossny. "Sentiment analysis over social Networks: An overview." In *2015 IEEE International Conference on Systems, Man, and Cybernetics*, pp. 2174-2179. IEEE, 2015.
 27. Carl-Gustav Linden. "The bright future of news automation." In *CEUR Workshop Proceedings*. 2018.
 28. Mahmood Yousefi-Azar and Len Hamey. "Text summarization using unsupervised deep learning." *Expert Systems with Applications* 68 (2017): 93-105. Jan A. Botha,

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

Manaal Faruqui, John Alex, Jason Baldrige, and Dipanjan Das. “Learning To Split and Rephrase From Wikipedia Edit History.” *arXiv [cs.CL]*. arXiv 2018.

<http://arxiv.org/abs/1808.09468>.

29. Perri 6. Don't Try this at Home: Lessons from England. Chapter 12 (pp. 325-354) in Sanford Borins, Kenneth Kernaghan, David Brown, Nick Brontis, Perri 6 and Fred Thompson (eds.) *Digital State at the Leading Edge*. Toronto: University of Toronto Press, 2007.
30. One approach to evaluating the viability of an AIPA-generated briefing note as being adequate for some purposes is to conduct A/B experiments with either real or proxy decision makers using briefing notes generated using AIPA, and parallel notes from human analysts.
31. A prominent example is the move from high-occupancy vehicle (HOV) lanes to high-occupancy smart-toll (HOST) lanes using real-time traffic and congestion data to dynamically adjust the price for drivers who do not meet the existing HOV conditions to use the lane. A pilot project in Los Angeles uses onboard occupancy transponders, data from the toll and traffic systems, parking data, public transit ridership data (riders earn toll charge credits), a violation processing system, traffic cameras, and dynamic messaging signs to update price recommendations every 5 minutes. Drivers see the current price and decide whether to enter the lane based on their own criteria. See Justin Longo. 2018. *Digital Tools for Rapid Policy Design*. In M. Howlett and I. Mukherjee (eds.). *Handbook of Policy Design*. Chapter 19, pp. 288-301. New York: Routledge.
32. Justin Longo and Kathy McNutt. “From Policy Analysis to Policy Analytics.” In *Policy Analysis in Canada*, edited by Laurent Dobuzinskis and Michael Howlett. International Library of Policy Analysis. Bristol UK: Policy Press, 2018.
33. Safiya Umoja Noble. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York University Press.

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

34. Tarleton Gillespie. 2014. "The Relevance of Algorithms." In *Media Technologies: Essays on Communication, Materiality, and Society*. Eds. T. Gillespie, P. Boczkowski, and K. Foot. MIT Press. 167-193.
35. Natascha Just and Michael Latzer. 2017. "Governance by Algorithms: Reality Construction by Algorithmic Selection on the Internet." *Media Culture & Society* 39 (2): 238-58.
36. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. "BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding." *arXiv [cs.CL]*. arXiv. 2018. <http://arxiv.org/abs/1810.04805>.
37. Rajpurkar, P., Zhang, J., Lopyrev, K., & Liang, P. (2016). SQuAD: 100,000+ Questions for Machine Comprehension of Text. *EMNLP*.
38. Douglas W. Oard, Jason R. Baron, Bruce Hedin, David D. Lewis, and Stephen Tomlinson. "Evaluation of information retrieval for E-discovery." *Artificial Intelligence and Law* 18, no. 4 (2010): 347-386.
39. See <https://openai.com/blog/better-language-models/>
40. While medical randomized control trials (RCTs) have developed good guidelines around informed consent and the ethics of placebos, deception and blinding (Edwards, S. J., R. J. Lilford, D. A. Braunholtz, J. C. Jackson, J. Hewison, and J. Thornton. 1998. "Ethical Issues in the Design and Conduct of Randomised Controlled Trials." *Health Technology Assessment* 2 (15): i - vi.), the politics of policy experimentation in a democracy raise legal and ethical questions about the unequal treatment of citizens in a true RCT (Pearce, Warren, and Sujatha Raman. 2014. "The New Randomised Controlled Trials (RCT) Movement in Public Policy: Challenges of Epistemic Governance." *Policy Sciences* 47 (4). Springer US: 387-402). Experiments designed to test different social welfare interventions—e.g., the Rand health care experiment, and various universal basic income experiments—have been attempted at small scale. However questions about fairness in a democratic society where a

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

new policy intervention is tested on some but not others, whether providing a benefit or imposing a cost, need to be considered.

41. A modification of the concept of clinical equipoise—a state of not knowing which intervention is best, or even whether any intervention is better than doing nothing—to one of ‘social equipoise’ in policy experiments, has been proposed. See Petticrew, M., M. McKee, K. Lock, J. Green, and G. Phillips. 2013. “In Search of Social Equipoise.” *BMJ* 347 (July): f4016.
42. Howard Rheingold. 1985. *Tools for Thought: The History and Future of Mind-expanding Technology*. MIT Press.
43. Hawksworth, J., Berriman, R., & Goel, S. (2018). Will robots really steal our jobs?: An international analysis of the potential long term impact of automation. Retrieved from Pricewaterhousecoopers website:
https://www.pwc.com/hu/hu/kiadvanyok/assets/pdf/impact_of_automation_on_jobs.pdf
44. Howard Rheingold. *Tools for Thought: The History and Future of Mind-expanding Technology*. MIT Press, 1985.
45. Douglas C. Engelbart. Augmenting human intellect: A conceptual framework. *Stanford Research Institute*, 49(638), 1024, 1962.
46. Canada. (2018). Digital Operations Strategic Plan: 2018-2022. Retrieved from Treasury Board of Canada Secretariat website:
<https://www.canada.ca/en/government/system/digital-government/digital-operations-strategic-plan-2018-2022.html>
47. Ling, J. (2018, May 31). Federal governments looks to AI in addressing issues with immigration system. *The Globe and Mail*. Retrieved from
<https://www.theglobeandmail.com/politics/article-federal-government-looks-to-ai-in-addressing-issues-with-immigration/>
48. Ling, 2018, p. np.

by: Ben Armstrong, Megan Beretta, Evan Crothers, Michael Karlin, Dongwoo Kim, Justin Longo, Lorne Powell and Trooper Sanders

49. Longo, J., Kuras, E., Smith, H., Hondula, D. M. & Johnston, E. (2017), Technology Use, Exposure to Natural Hazards, and Being Digitally Invisible: Implications for Policy Analytics. *Policy & Internet*, Vol. 9, no. 1, pp. 76-108 doi:10.1002/poi3.144.