Appendix:

Analysis Method: Bayesian Gaussian Process Regression

This *Commentary*'s aim is to estimate the likelihood that a given occupation can be automated, based on its skill content. The O*Net database used contains detailed information for 954 occupations, including the level of each skill required and importance of that skill to the job, as well as the level of various activities required.

The conceptual basis for the model is that there are some skills or tasks where humans still dominate robots. Jobs that require high levels of these skills cannot be fully automated, as humans are still better than machines at completing the job tasks. Conversely, jobs that do not require these skills, or where the skill is not important to the job, are more likely to be automatable.

The skill variables described in Table 1 fully specify an explanatory variable set or a *feature vector* for each occupation. The dependent variable (whether or not a job is automatable) is incomplete and, for the most part, unknown. I only know which jobs are fully automatable and, with some certainty, the jobs that are currently impossible to automate. For those jobs that are partially automatable, the level is unknown.

To handle this difficulty with the data, I employ a Bayesian Gaussian process regression, which is a powerful kernel-based classification method. It is a method that uses "training data" – the information about which I am certain – to probabilistically determine the remaining unknown values. This method also does not limit the relationship between skills and the potential for automation to constant relationships (ie. how much a certain skill effects the chances of being automated is not constant over all automation levels).

There are some important limitations to this method of analysis. The results should be interpreted as the probability that an occupation could be automated in theory, or the approximate percentage of a given occupation that could be automated. It does not account for growth in demand as a result of labour-productivity improvements nor the creation of new occupations as a result of technological change. Despite these limitations, analyses of automation in the Canadian labour market have used some variant of a similar regression analysis, allowing for the current analysis to be compared to previous results.

A Gaussian process is defined as a collection of random variables, any finite number of random variables that have a joint Gaussian distribution (Rasmussen and Williams 2006).[27] Gaussian processes have been extensively used for many different variants of nonparametric estimation (Seeger 2002). In addition, this method has been applied to estimate occupations' susceptibility to automation in previous studies (Frey and Osborne 2013, Oschinski and Wyonch 2017).

A classification of some occupations as fully automatable ($y_i = 1$), not automatable ($y_i = 0$) or partially automatable ($y_i = 0.5$) is required to "train" the model. To develop a classification, I use the output classification results of previous studies on occupational automation: Autor and Dorn (2013), Josten and Lorden (2019), Frey and Osborne (2013) and Oschinski and Wyonch (2017). With the exception of Autor and Dorn (2013), which used binary classification, these studies classified occupations into three categories. Across the classification outputs, there were many occupations that were given a similar classification. Those were

27 For an extensive discussion of the applications and theory of Gaussian processes for machine learning and statistical analysis, see Rasmussen & Williams (2006).

classified equivalent to previous research. Occupations that were classified as "medium risk" in the studies using trinary classification were classified similarly for the training set, ignoring the binary classification from Autor and Dorn (2013) in these cases. Where previous research did not agree on a particular classification, no class was designated in the training data. This effectively combines the information about occupational automation that has been consistent across different methods of analysis.

The training data set is defined by $\mathcal{D} = (X, \boldsymbol{y})$, where X is the matrix of feature vectors and $\boldsymbol{y}$ gives the associated class label. Each element $x_{i,j}$ of X represents the level of each skill (j) required, weighted by the importance of that skill to the specified occupation (i). I assume $\mathcal{D} = (X, \boldsymbol{y})$, $\boldsymbol{x}_i \in \mathrm{R}^9$, $y_i \in \{0,1\}$ is a noisy independent, identically distributed sample from latent function $f : \boldsymbol{x} \rightarrow R$ where $\boldsymbol{w} = P(y \mid f)$ where denotes the noise distribution.

Given dataset of observations, $\mathcal{D} = \{(\boldsymbol{x}_i, y_i) \mid i = 1,2, \ldots, n\}$, we wish to make predictions ($\boldsymbol{y}_*$) for new input features X_* that are not in the training set $\mathcal{D}$. To estimate $\boldsymbol{y}_*$, I employ a zero-mean Gaussian process prior ($\boldsymbol{w} \sim N(0,\Sigma)$) and a generative Bayesian method. The process gives a prior probability weight to every possible function that describes the relationship specified by $\mathcal{D}$, where higher probabilities are assigned to functions I consider more likely. The likelihood of a function is determined by its relative proximity to training data points and the prior specification that fixes the properties of functions to be considered for inference.

The model is then computed by:

(1) Introduce $\phi(\boldsymbol{x})$ which maps input vector $\boldsymbol{x}$ into an N-dimensional feature space. Further, let $\Phi(X)$ be the aggregation of the columns $\phi(\boldsymbol{x})$. The model is defined by:

$$f(\boldsymbol{x}) = \phi(\boldsymbol{x})^{\mathrm{T}} \boldsymbol{w}$$

(2) Conditioning the prior distribution on the data $\mathcal{D}$, resulting in a posterior distribution:

$$posterior = \frac{likelihood \times prior}{marginal\ likelihood}, \qquad p(w|\boldsymbol{y}, \Phi(X)) = \frac{p(\boldsymbol{y}|\Phi(X), \boldsymbol{w})p(w)}{p(y|\Phi(X))}$$

$$p(\boldsymbol{w}|\boldsymbol{y}, \Phi(X)) \sim N\left(\bar{w} = \frac{1}{\sigma_n^2} A^{-1}\Phi(X)\boldsymbol{y}, A^{-1}\right)$$

Where $A = \sigma n^{-2}\, \Phi(X)\Phi(X)^{\mathrm{T}} + \Sigma^{-1}$

(3) The marginal posterior is the predictive distribution.

$$f_*|\boldsymbol{x}_*, \mathcal{D} = \int p(f_*|\phi_*, w)p(w|\Phi, \boldsymbol{y})dw \ \sim N(\frac{1}{\sigma_n^2}\phi_*{}^T A^{-1}\Phi\boldsymbol{y}, \phi_*{}^T A^{-1}\phi_*)$$

Where $\Phi = \Phi(X)$ and $\phi^* = \phi(\boldsymbol{x}_*)$

An alternative formulation is given by:

$$f_*|\boldsymbol{x}_*, \mathcal{D} \sim N(\phi_*^T \Sigma \Phi (K + \sigma_n^2 I)^{-1} \boldsymbol{y},\ \ \phi_*^T \Sigma \phi_* - \phi_*^T \Sigma \Phi (K + \sigma_n^2 I)^{-1} \Phi^T \Sigma \phi_*) \tag{5}$$

Where $K = \Phi^T \Sigma \Phi$

Notice that in equation (5) the feature space always enters in a form that dictates entries in the matrices are of the form $\phi(\boldsymbol{x})^T \Sigma \phi(\boldsymbol{x}')$ where $\boldsymbol{x}$ and $\boldsymbol{x}'$ are in either the training set or the test set.

Let $k(x,x') = \phi(x)^T \Sigma \phi(x')$ be the kernel or covariance function. The specification of a covariance function implies a distribution over functions. A Gaussian process is completely specified by its mean function and covariance function.

$$f(\boldsymbol{x}) \sim GP\ (m(\boldsymbol{x}),\ \mathrm{k}(\boldsymbol{x},\boldsymbol{x}')) \tag{6}$$

$$\boldsymbol{y} = f(\boldsymbol{x}) + \varepsilon \tag{7}$$

For our model $f(\boldsymbol{x}) = \phi(\boldsymbol{x})^T \boldsymbol{w}$ with prior $\boldsymbol{w} \sim N(\boldsymbol{0},\Sigma)$ I have mean and covariance

$$E[f(\boldsymbol{x})] = 0 \tag{8}$$

$$E[f(\boldsymbol{x})f(\boldsymbol{x}')] = (\boldsymbol{x})^T \Sigma \phi(\boldsymbol{x}') = k(\boldsymbol{x}, \boldsymbol{x}') \tag{9}$$

Notice, the covariance between outputs is a function of inputs. I define the kernel as the squared exponential covariance function:

$$k(\boldsymbol{x}, \boldsymbol{x}') = \exp\left(-\frac{1}{2}|\boldsymbol{x} - \boldsymbol{x}'|^2\right) \tag{10}$$

Assuming additive, identically distributed Gaussian noise with variance , the prior on noisy observations becomes:

$$cov(\boldsymbol{y}) = K(X,X) + \sigma_n^2 I. \tag{11}$$

The free parameter σ_n^2 is referred to as a hyperparameter to emphasize that it refers to a parameter of a non-parametric model. The parameters (weights) of the underlying parametric model have been integrated out.

To compute the model, I employ the "kernlab" R package (Karatzoglou, Smola and Hornik 2019). Specifically, I use the "gausspr" function, an exponentiated quadratic covariance (radial basis function kernel = "rbfdot") and an optimized hyperparameter selection. The accuracy of the model is estimated by a 10-fold cross validation error, which yields an error estimate of 0.09 on predicted probabilities.

Attribute Selection and Sensitivity Analysis

The attributes selected as barriers to automation in the main analysis are detailed in Table 1. There are, however, additional factors that could affect the likelihood that an occupation could be automated (Table A1). Similarly, results will be affected by the classification given for occupations in the training set.

Table A1: Attributes that Pose Barriers to Automation

Social Perception	Being aware of others' reactions and understanding why they react as they do.
Originality	The ability to come up with unusual or clever ideas about a given topic or situation, or to develop creative ways to solve a problem.
Assisting and Helping Others	Providing personal assistance, medical attention, emotional support or other personal care to others such as coworkers, customers, or patients.
Philosophy	Knowledge of different philosophical systems and religions. This includes their basic principles, values, ethics, ways of thinking, customs, practices and their impact on human culture.
Initiative	Willingness to take on responsibilities and challenges.
Leadership	Willingness to lead, take charge, and offer opinions and direction.
Innovation	Creativity and alternative thinking to develop new ideas for and answers to work-related problems.
Adaptability and Flexibility	Being open to change (positive or negative) and to consider variety in the workplace.
Independence	Developing one's own way of doing things, guiding oneself with little or no supervision and depending on oneself to get things done.
Complex Problem Solving	Identifying complex problems and reviewing related information to develop and evaluate options and implement solutions.
Resolving Conflict and Negotiating with Others	Handling complaints, settling disputes and resolving grievances and conflicts, or otherwise negotiating with others.
Repairing or Maintaining Electronic or Mechanical Equipment	Servicing, repairing, calibrating, regulating, fine-tuning or testing machines, devices, equipment, and moving parts.
Public Speaking or Interpreting the Meaning of Information for Others	Frequency of public speaking, translating or explaining what information means and how it can be used.
Developing Objective Strategies	Establishing long-range objectives and specifying the strategies and actions to achieve them.

Source: O*NET Database.

Table A2: Models Specifying Different Attributes Posing Barriers to Automation

	Main analysis	Expanded attribute selection	Work activities only
Attributes	Social perception, originality, assisting and helping others, knowledge of philosophy, initiative, leadership, innovation, independence, adaptability and flexibility	Social perception, originality, assisting and helping others, initiative, leadership, innovation, independence, adaptability and flexibility, complex problem solving, resolving conflict and negotiating with others, repairing and maintaining equipment, public speaking and developing objective strategies	Assisting and helping others, complex problem solving, resolving conflict and negotiating with others, repairing and maintaining equipment, and developing objective strategies
Mean Automation Risk Across Occupations (unweighted %)	48.4	45.7	46.1
Correlation of Estimated Probability of Automation (Model output)			
Main Analysis	1	0.954	0.889
Expanded Attribute Selection	0.954	1	0.921
Work Activities only	0.889	0.921	1

Source: Author's calculations.

Table A3: Models Specifying Different Input Classifications – Correlation and Average Risk of Automation

	Oschinski and Wyonch	Frey and Osborne	Meta	Autor and Dorn	Josten and Lorden
Oschinski and Wyonch	1	0.93659	0.943999	0.636234	0.822301
Frey and Osborne		1	0.93865	0.642068	0.803695
meta			1	0.828096	0.941843
Autor and Dorn				1	0.861623
Josten and Lorden					1
Unweighted average automation risk	0.56	0.54	0.48	0.40	0.41

Source: Author's calculations.

Table A4: Change in Employment by Risk Category (2020 to 2028)				
Technology Adoption Rate (%, annual)	2.1	1.7	4.3	7.5
Number of Employees				
Low Risk	490,000	522,000	254,000	-58,000
Med Risk	-262,000	-176,000	-878,000	-1,611,000
High Risk	-322,000	-260,000	-753,000	-1,236,000
Net Employment	-93,000	87,000	-1,377,000	-2,906,000
Percent Change				
Low Risk	3.3	3.4	1.8	-0.1
Medium Risk	-1.4	-0.9	-5.2	-9.7
High Risk	-1.9	-1.5	-4.5	-7.5

Source: Author's calculations.

The attributes listed as barriers to automation are those identified by similar research studies included in the analysis. A selection of these attributes is included in the main analysis. Alternative analysis using an expanded selection and one that includes work activities only (similar to Autor and Dorn 2013) shows that addition or removal of individual attributes has little effect on aggregate results (Table A2). It does, however, have marginal effects on the classification of particular occupations as low, medium or high risk of automation.

Skills or attributes that enable humans to work with technology but could also be performed by computers or machines, such as data analysis and pattern recognition, are ambiguous in their impact on automation. For attributes or skills to pose a barrier to automation, by definition, machines/computers cannot outperform humans at them. Analytical skills or digital competencies are likely to be growing aspects of many, if not all occupations as technology continues to be developed and adopted. A recent report from the UN Industrial Development Organization summarizes the relationship between human skills and automating technologies (p. 75):

> "Technological change tends to favour skills that are complementary to the new technology (Acemoglu 2002, Rodrik 2018). Even if debate on the set of skills that will be required to perform with ADP technologies is still open, the needed skills are expected to be biased towards three broad categories: analytical, technology-related and soft skills (Kupfer et al. 2019)."

For analytical and technology-related skills, humans might compliment machines, but they might also compete with them – so while they are complimentary to technology, having advanced analytical skills does not necessarily mean that the analysis couldn't be automated. Those who have strong technology-related and analysis skills might be better able to adapt to changing business practices but many analysis-based tasks can be significantly automated. Improving these skills throughout the population regardless of age, income or occupation would likely reduce the negative employment effects of technology adoption and growth. In addition, a highly skilled workforce enables businesses to more readily adopt productivity improving technologies that increase output without changing the demand for labour.

This analysis focusses on the risk of an occupation being automated, and not on what efficiency improvements technology adoption might lead to or how it will be complimented by human labour. The method used to calculate the likelihood of automation requires that the skills and attributes included in the model be exclusively human, or at least areas where humans are likely to outperform machines for quite some time.

To determine whether the training input classification significantly affects the results of the analysis, I also estimated the probability that an occupation is automatable using classifications from Autor and Dorn (2013), Lorden and Jorden (2019), Frey and Osborne (2013) and Oschinski and Wyonch (2017) independently. Similarly, to determine whether the attribute selection significantly affects the results, I estimated the model with different combinations of attributes and a similar classification input.

Collectively, the results are quite consistent across differing input classifications and explanatory attribute selections. Varying the selection of attributes that pose barriers to automation yields average likelihoods of automation of about 46 percent to 48 percent (Table A2). In addition, the estimated likelihood of automation is highly correlated across the variant selections tested here. Similarly, probability estimates calculated using differing input classifications resulted in broadly similar results (Table A3). There is, however, significant variability in the estimates for individual occupations. Across the seven model selections specified in Tables A2 and A3, the average gap in estimated probability of automation for individual occupations is 30 percent (maximum gap is 0.82, minimum gap is 0.06).

The estimates associated with the combined classification (labelled "meta" in Table A3) are highly correlated with those calculated from each classification individually. The average likelihood of automation falls in the middle of the estimates. Similarly, the probability of automation estimates for individual occupations all fall between the extremes. Due to remaining uncertainty about which classification most reflects real-world potential for automation and the rather convenient result that the combined classification yields estimates that fall between the extremes estimated in other models, the "meta" classification is used for the main analysis of the potential effects of automation on Canada's labour market.

Mapping O*NET data to Canadian Labour Market Data

The statistical analysis estimates the likelihood an occupation can be automated based on the selected attributes and the classifications in the training vector. The skills and attribute information for each occupation is sourced from the O*NET database, sponsored by the US Department of Labor. O*NET data are based on ratings from employers, workers and occupational analysts. The O*NET data are coded to the 2010 Standard Occupational Classification (SOC2010). To apply the O*NET information to Canadian occupational data, I utilize a concordance developed by Statistics Canada and follow the method used in Frenette and Frank (2017):

"A concordance between the six-digit SOC2010 codes (provided in the O*NET database) and the four-digit 2011 National Occupational Classification (NOC2011) codes was developed. This concordance was determined based on the similarity of the occupational descriptions.

In some instances, a six-digit SOC2010 code had "breakout" (i.e., new or emerging) sub-occupations. Since there were no estimates of population size for these sub-occupations, it was assumed that they were of equal size within a six-digit SOC2010 code. After the breakout occupations were aggregated, 1,058 SOC2010–NOC2011 matched pairs remained. This number included 495 unique four-digit NOC2011 codes.

Of these 495 unique NOC2011 codes, 336 were matched with a six-digit SOC2010 code. For the remaining 159 NOC2011 codes, higher-level SOC2010 codes had to be used, because more than one six-digit SOC2010 code mapped to them (i.e., the ACS data were used at this stage). Among these 159 NOC2011 codes, 119 were matched with five-digit SOC2010 codes, 38 were matched with four-digit SOC2010 codes and the remaining two were matched with three-digit SOC2010 codes.

At this stage, the 495 unique NOC2011 codes contained the information on occupational skill requirements from O*NET. The data on occupational skill requirements were then linked to the NOC2011 codes."