Machine Learning in Healthcare? Evidence from online job postings

By Avi Goldfarb, Bledi Taska, and Florenta Teodoridis

* Goldfarb: Rotman School of Management, University of Toronto, 105 St George St, Toronto, ON M5M1N8 (avi.goldfarb@rotman.utoronto.ca), Taska: Burning Glass Technologies, One Lewis Wharf, Boston, MA 02110, (btaska@burning-glass.com), Teodoridis: University of Southern California, Marshall School of Business, 701 Exposition Blvd., Los Angeles, CA 90089 (florenta.teodoridis@marshall.usc.edu). Thanks to seminar participants at the University of Toronto, Harvard University and the NBER for helpful comments.

Since the 1970s, computer scientists and medical professionals have recognized the potential for artificial intelligence (AI) to transform healthcare (Shortliffe 1976, Clancey and Shortliffe 1984, Fieschi 1990). Recent advances in machine learning have brought renewed attention to AI generally (Agrawal, Gans, and Goldfarb 2018), and AI in medicine in particular (Mukherjee 2017; Topol 2019).

This paper documents a puzzle. Despite the numerous popular press discussions of AI in healthcare, there has been relatively little adoption. Using data from Burning Glass Technologies on millions of online job postings over ten years, we find that AI adoption in healthcare remains substantially less than in most other industries.

Figure 1 shows the fraction of jobs requiring AI skills, by industry (defined by 2-digit NAICS). The healthcare and social assistance category has the second lowest rate of AI jobs of the 20 industries in the data. This low rate of adoption is not driven by the combination of healthcare with social assistance. Roughly 1 in 1,250 hospital jobs required AI-related skills in 2015-2018 compared to approximately 1 in 174 in finance & insurance, 1 in 88 in professional, scientific, and technical services, and 1 in 72 in information. This low rate of adoption in healthcare is the main puzzle identified in this paper.[1]

I. Data

Our focal data set comes from Burning Glass Technologies and is based on over 40,000 online job boards and company websites.[2] It includes 93,237,194 job postings in the United States from January 1, 2015 to December 31,

[1] The low rate of adoption and early stage of the technology also means that our analysis is necessarily descriptive. We cannot reliably link AI adoption with patient outcomes or with hospital efficiency, a key focus on the prior literature on healthcare information technology (Agha 2014; Dranove et al 2014; Lee, McCollough, and Town 2014; McCollough, Parente, and Town 2016; Freedman, Lin, and Prince 2018; Lu, Rui, Seidmann 2018)

[2] The strengths and weaknesses of this dataset are detailed in Hershbein and Kahn (2018) who use it to measure the impact of recessions on skill requirements in job vacancies. This data set is representative of vacancies in healthcare and many other industry sectors. It over-samples professional and skilled jobs relative to survey-based measures of vacancies.

2018. We combine this with American Hospital Association (AHA) data from 2014 on over 200 hospital characteristics of 4,556 US hospitals. We identify 1,840,784 job postings for these hospitals in the Burning Glass data between 2015 and 2018. In addition, we use data from County Business Patterns in 2013 to examine the role of location-specific characteristics in adoption, similar to Forman, Goldfarb, and Greenstein (2008).[3]

To define adoption of AI, we leverage an insight in Tambe and Hitt (2012): That hiring decisions can be used to understand technology adoption and diffusion. For each job posting, Burning Glass classifies the skills listed into a number of skill clusters. We define AI as those jobs that they categorize into the following skill clusters: "Artificial Intelligence", "Machine Learning", and "Natural Language Processing". Examples of AI jobs in our data include "Post-Doctoral Fellow in Cardiovascular Genomics and Biomedical Informatics" (listing machine learning as a required skill), "Analytics Architect" (listing both machine learning and artificial intelligence), and "Population Management Educator" (listing IBM Watson). In total, we find 1,479 AI job postings at 126 different hospitals.[4]

In addition to identified which job postings require AI skills, we use the Burning Glass data to identify clinical, research, and administrative jobs (respectively 60%, 6%, and 34% of job postings).[5] Furthermore, from the AHA data, we identify whether hospitals use an integrated salary model to pay doctors. In an integrated salary model, doctors are given a salary from the hospital, rather than being paid for each service provided. 53% of the job postings are in hospitals with an integrated salary model. Finally, the AHA data has a measure of whether a hospital is a teaching hospital, defined as having a residency program or being a member of the council of teaching hospitals, along with other variables we use as controls. In the AHA data, some hospitals are missing important data, such as the integrated salary model variable and the

[3] Using principal components analysis, we identified the most important variables in the AHA and County Business Patterns data. In the AHA data, the most important component (with proportion 0.23) related to hospital size which we proxy with total full-time employees. For the County Business Patterns data, the most important component related to population (with proportion 0.82), which we proxy with county payroll. We purposefully select hospital and county measures from before our observation period (2015-2018) to avoid confounding effects of adoption influencing these metrics.

[4] Our qualitative results are robust to using a broader definition that includes the cluster "Data Sciences". With this broader definition, we identify 2,869 job postings from 233 hospitals.

[5] We define research job postings as posting where "research" is in the job title or which lists skills in any of the following skills clusters: "Clinical Research", "Laboratory Research", "Research Methodology", and "Research And Development Industry Knowledge". We define clinical jobs postings as non-research job postings classified under one of the following Standard Occupation Classification (SOC) codes: 29-000 "Healthcare practitioners and technical Occupations" or 31-0000 "Healthcare Support Occupations". We define administrative jobs postings as all other non-research, non-clinical job postings.

hospital size variable. Our analysis uses the 1,582,333 job postings from hospitals for which we have this data.

II. Empirical analysis

Our empirical analysis takes the job posting as the unit of observation.

A. What correlates with adoption?

Table 1 shows logit regressions of a dummy for whether the job posting lists AI skills on whether the hospital has an integrated salary model, county payroll, hospital size, and a variety of other controls.[6] Coefficients are shown with standard errors clustered by hospital in parentheses.

Consistent with prior work on adoption of information technology in healthcare and in other industries (e.g. Dranove et al 2014; Forman, Goldfarb, and Greenstein 2008), adoption is more likely in larger counties and in larger hospitals.

In addition, all three columns show that hospitals governed by an integrated salary model are more likely to post jobs that require AI skills. Column (2) shows that clinical jobs are somewhat less likely to require AI skills compared to research jobs and administrative jobs. In the raw data, there are just 105 clinical job postings from 34 hospitals. Although there are far more clinical postings than administrative ones, there are 965 administrative AI postings in 94 hospitals.

This contrasts with the discussion in the popular press that emphasizes clinical applications (Mukherjee 2017; Topol 2019); however, it is consistent with the relatively widespread use of prior generations of data science techniques in research and information technology tools in billing (Adams-Huet and Ahn 2009, Atasoy, Greenwood, and McCollough 2019).

B. Why might healthcare lag?

Any interpretation of the above results should recognize the very low overall adoption rate. In our view, the most important result in the data is shown in figure 1: hiring for AI skills in healthcare remained rare by the end of 2018.

There are a variety of reasons why healthcare might lag, including regulatory constraints, the early stage of the technology, managerial challenges related to coinvention (Bresnahan and Greenstein 1996), the practical usefulness of the technology in healthcare, data limitations (for example AI software used in hospitals but

[6] As previously mentioned, our set of controls is informed by a principal component analysis. We include surrogates for the top by explanatory power, in addition to the main explanatory variable, hospital size, which we proxy with total full-time employees. The surrogates are: preferred provider organization (PPO) hospital, general medical and surgical care (adult) hospital, psychiatric care hospital, acute long-term care hospital, nutrition program hospital, and rural health clinic hospital.

developed by non-healthcare companies), and misaligned incentives to adopt.

We provide some suggestive evidence of misaligned incentives, though we do not reject the possibility that the other factors also matter a great deal, and perhaps more.

We interpret our results to suggest potential for misaligned incentives for two reasons. First, column (3) shows that higher adoption under the integrated salary model is not true for research postings, in which AI might enhance the capabilities of the physician-scientist (Adams-Huet and Ahn 2009). Instead, it is driven by clinical and administrative postings, where the positions of those hiring might be threatened by the technology (Jamieson and Goldfarb 2019).[7]

Second, Columns (4) and (5) show that teaching hospitals are not more likely to hire for AI positions. The variable identifying teaching hospitals provides a different type of measure of hospital quality in the AHA data. Teaching hospitals are often seen as better-quality hospitals (Ayanian and Weissman 2002).

Hospitals with an integrated salary model differ from other hospitals in multiple dimensions such as the degree of professional management and the allocation of decision rights.

While the variable identifying teaching hospitals may only loosely correlate with management quality, we see the identical coefficients in columns (1) and (5) on the integrated salary model variable as suggestive that the integrative salary result is not driven by a vertical measure of managerial quality. The integrated salary model variable appears to be capturing something else. Combined with the results on research adoption, we think incentives to adopt technology that replaces hospital decision-makers is a possible explanation.[8]

It is important to emphasize that this analysis is suggestive. While our descriptive statistics are consistent with one barrier to adoption being incentives to avoid job displacement, it does not carefully analyze possible omitted variables.

III. Conclusions

Using data on the skills requires in online job postings to measure AI adoption, we showed that AI adoption in healthcare generally and in hospitals in particular, is low. Approximately 1

[7] This is consistent with more causally-focused research showing the importance of incentives in medical treatment (Gruber, Kim, and Mayzlin 1999; Clemens and Gottleib 2014).

[8] Several studies have explored the vulnerability of occupations to AI and automation (e.g. Brynjolfsson, Mitchell, and Rock 2018; Felten, Raj, and Seamans 2018). We think using these vulnerability scores to test the hypothesis of misaligned incentives is a promising avenue of research; however, it requires data on the occupation of the person deciding to make a hire rather than the data we have on the occupation of the person hired.

in 1,250 hospital jobs posted between 2015 and 2018 required AI skills, and under 3% of the 4,556 hospitals in our data posted any jobs requiring AI skills over this time period. The low adoption rates mean any statistical analysis is limited. Nevertheless, the adoption we do see in the data shows that larger hospitals, larger counties, and integrated salary model hospitals are more likely to adopt.

REFERENCES

Adams-Huet, Beverley, and Chul Ahn. 2009. "Bridging clinical investigators and statisticians: writing the statistical methodology for a research proposal." *Journal of investigative medicine* 57(8), 818-24.

Agha, Leila. 2014. The effects of health information technology on the costs and quality of medical care. *Journal of Health Economics* 34, 19-30.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018. *Prediction Machines*. Harvard Business Review Press, Boston MA.

Atasoy, Hilal, Brad N. Greenwood, and Jeffrey Scott McCullough. 2019. The Digitization of Patient Care. *Annual Review of Public Health* 40, 487-500.

Ayanian, John Z, and Joel S Weissman. 2002. "Teaching hospitals and quality of care: a review of the literature." *The Milbank quarterly* 80(3), 569-93.

Bresnahan, Tim and Shane Greenstein. 1996. Technical Progress and Co-invention in Computing and in the Uses of Computers. *Brookings Papers on Economic Activity, Microeconomics* 1996; 1-83.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings*, 108, 43-47.

Clancey, William, and Edward H. Shortliffe, eds. 1984. *Readings in Medical Artificial Intelligence: The First Decade*. Addison Wesley, Reading MA.

Clemens, Jeffrey, and Joshua D. Gottlieb. 2014. "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?" *American Economic Review*, 104(4), 1320-49.

Dranove, David, Chris Forman, Avi Goldfarb, and Shane Greenstein. 2014. The Trillion Dollar Conundrum: Complementarities and Health Information Technology. *American Economic Journal: Economic Policy* 6(4), 239-270.

Felten, Edward W., Manav Raj, and Robert Seamans. 2018. "A Method to Link Advances in Artificial Intelligence to

Occupational Abilities." *AEA Papers and Proceedings*, 108, 54-57.

Fieschi, M. 1990. “Artificial Intelligence in Medicine: Expert Systems” Chapman & Hall

Forman, Chris, Avi Goldfarb, and Shane Greenstein. 2008. “Understanding Inputs into Innovation: Do Cities Substitute for Internal Firm Resources?” *Journal of Economics and Management Strategy* 17(2), 295-316.

Freedman, Seth, Haizhen Lin, and Jeffrey Prince. 2018. “Information Technology and Patient Health.” *American Journal of Health Economics* 4(1), 51-79.

Gruber, Jonathan, John Kim, and Dina Mayzlin, “Physician Fees and Procedure Intensity: The Case of Cesarean Delivery,” Journal of Health Economics, 1999, 18 (4), 473–490.

Hershbein, Brad, and Lisa B. Kahn. 2018. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review* 108 (7): 1737-72.

Jamieson, Trevor, and Avi Goldfarb. 2019. Clinical considerations when applying machine learning to decision-support tasks versus automation. *BMJ Quality & Safety* 28, 778-781.

Lee, Jinhyung, Jeffrey McCullough, and Robert Town. 2013. “The impact of health information technology on hospital productivity.” *RAND Journal of Economics*, 44, 545-568.

Lu, Susan, Huaxia Rui, and Abraham Seidmann. 2018. Does Technology Substitute for Nurses? *Management Science* 64(4), 1842-1859.

McCullough, Jeffrey, Stephen Parente, and Robert Town. 2016. “Health information technology and patient outcomes: the role of information and labor coordination.” *RAND Journal of Economics* 47, 207-236.

Mukherjee, Siddhartha. 2017. A.I. Versus M.D. *The New Yorker*. April 3.

Shortliffe, Edward Hance. 1976. *Computer-Based Medical Consultations: MYCIN*. Elsevier, New York.

Tambe, Prasanna, and Lorin Hitt. 2012. The Productivity of Information Technology Investments: New Evidence from IT Labor Data. *Information Systems Research* 23(3), 599-617.

Topol, Eric. 2019. *Deep Medicine*. Basic Books, New York.

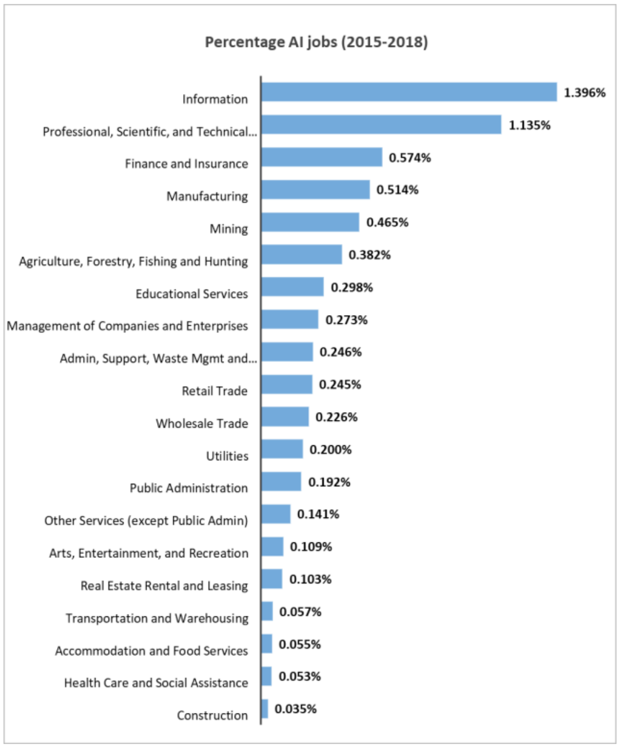

FIGURE 1. FRACTION OF JOB POSTINGS BY YEAR THAT LIST A.I. SKILLS AS A JOB REQUIREMENT

TABLE 1: LOGIT REGRESSION OF AI SKILLS ON HOSPITAL AND JOB CHARACTERISTICS

	(1)	(2)	(3)	(4)	(5)
Integrated Salary	0.840	0.764	1.124		0.840
	(0.251)	(0.226)	(0.267)		(0.251)
Teaching hospital				-0.080	-0.079
				(0.516)	(0.552)
Research Posting		0.243	1.105		
		(0.244)	(0.201)		
Clinical Posting		2.528	-2.022		
		(0.318)	(0.332)		
Integrated Salary			-1.108		
Research Posting			(0.321)		
Integrated Salary			-0.655		
Clinical Posting			(0.541)		
Log county	0.486	0.388	0.392	0.407	0.489
(millions)	(0.175)	(0.158)	(0.157)	(0.172)	(0.180)
Log full time	0.679	0.581	0.588	0.784	0.690
hospital	(0.210)	(0.213)	(0.212)	(0.249)	(0.251)
Pseudo-R^2	0.109	0.164	0.167	0.101	0.109

Unit of observation is the job posting. Dependent variable is AI skills mentioned in the job posting. All models are logit; standard errors in parentheses are clustered at the hospital level. Data includes all job postings by hospitals in the AHA data from 2015-2018 (1,582,333 observations). Regressions include year fixed effects and hospital-level controls.