

## ARTICLE OPEN

## Medical device surveillance with electronic health records

Alison Callahan<sup>1</sup>, Jason A. Fries<sup>1,2</sup>, Christopher Ré<sup>2</sup>, James I. Huddleston III<sup>3</sup>, Nicholas J. Giori<sup>3,4</sup>, Scott Delp<sup>5</sup> and Nigam H. Shah<sup>1</sup>

Post-market medical device surveillance is a challenge facing manufacturers, regulatory agencies, and health care providers. Electronic health records are valuable sources of real-world evidence for assessing device safety and tracking device-related patient outcomes over time. However, distilling this evidence remains challenging, as information is fractured across clinical notes and structured records. Modern machine learning methods for machine reading promise to unlock increasingly complex information from text, but face barriers due to their reliance on large and expensive hand-labeled training sets. To address these challenges, we developed and validated state-of-the-art deep learning methods that identify patient outcomes from clinical notes without requiring hand-labeled training data. Using hip replacements—one of the most common implantable devices—as a test case, our methods accurately extracted implant details and reports of complications and pain from electronic health records with up to 96.3% precision, 98.5% recall, and 97.4% F1, improved classification performance by 12.8–53.9% over rule-based methods, and detected over six times as many complication events compared to using structured data alone. Using these additional events to assess complication-free survivorship of different implant systems, we found significant variation between implants, including for risk of revision surgery, which could not be detected using coded data alone. Patients with revision surgeries had more hip pain mentions in the post-hip replacement, pre-revision period compared to patients with no evidence of revision surgery (mean hip pain mentions 4.97 vs. 3.23;  $t = 5.14$ ;  $p < 0.001$ ). Some implant models were associated with higher or lower rates of hip pain mentions. Our methods complement existing surveillance mechanisms by requiring orders of magnitude less hand-labeled training data, offering a scalable solution for national medical device surveillance using electronic health records.

npj Digital Medicine (2019)2:94

; <https://doi.org/10.1038/s41746-019-0168-z>

## INTRODUCTION

For individuals implanted with the most widely used medical devices today, including pacemakers, joint replacements, breast implants, and more modern devices, such as insulin pumps and spinal cord stimulators, effective pre-market assessment and post-market surveillance are necessary<sup>1</sup> to assess their implants' safety and efficacy. Device surveillance in the United States relies primarily on spontaneous reporting systems as a means to document adverse events reported by patients and providers, as well as mandatory reporting from manufacturers and physicians to the Food and Drug Administration (FDA). As a result, device-related adverse events are significantly underreported—by some estimates as little as 0.5% of adverse event reports received by the FDA concern medical devices.<sup>2</sup> Recent adoption of universal device identifiers<sup>3,4</sup> will make it easier to capture device details, and to report events using these identifiers, but linking devices to patient outcomes remains a challenge. Unique device identifiers have not been used historically (longitudinal data are essential to monitor device safety over time) and many relevant patient outcomes are simply not recorded in a structured form amenable for analysis. The International Consortium of Investigative Journalists recently analyzed practices of medical device regulation around the world and dissemination of device recalls to the public, sharing their findings in a report dubbed the 'Implant Files'.<sup>5</sup> One issue highlighted in their report is the significant

challenge of tracking device performance and reports of adverse events.

Observational patient data captured in electronic health records (EHRs) are increasingly used as a source of real-world evidence to guide clinical decision making, ascertain patient outcomes, and assess care efficacy.<sup>6–8</sup> Clinical notes in particular are a valuable source of real-world evidence because they capture many complexities of patient encounters and outcomes that are underreported or absent in billing codes. Prior work has demonstrated the value of EHRs, and clinical notes specifically, in increasing the accuracy<sup>9</sup> and lead time<sup>10</sup> of signal detection for post-market drug safety surveillance. The hypothesis motivating this work is that EHRs, and evidence extracted from clinical notes in particular, can enable device surveillance and complement existing sources of evidence used to evaluate device performance.

EHR structure and content varies across institutions, and technical challenges in using natural language processing to mine text data<sup>11</sup> create significant barriers to their secondary use. Pattern and rule-based approaches are brittle and fail to capture many complex medical concepts. Deep learning methods for machine reading in clinical notes<sup>12–15</sup> outperform these approaches to substantially increase event detection rates, but require large, hand-labeled training sets.<sup>16,17</sup> Thus recent research has led to new methods that use weak supervision, in the form of heuristics encoding domain insights, to generate large amounts of

<sup>1</sup>Center for Biomedical Informatics Research, Stanford University, 1265 Welch Road, Stanford, CA, USA94305; <sup>2</sup>Department of Computer Science, Stanford University, 353 Serra Mall, Stanford, CA, USA94305; <sup>3</sup>Department of Orthopaedic Surgery, School of Medicine, Stanford University, 450 Broadway Street, Redwood City, CA, USA94063; <sup>4</sup>Veterans Affairs Palo Alto Health Care System, 3801 Miranda Avenue, Palo Alto, CA, USA94304 and <sup>5</sup>Department of Bioengineering, Stanford University, 318 Campus Drive, Stanford, CA, USA94305

Correspondence: Alison Callahan (acallaha@stanford.edu)

These authors contributed equally: Alison Callahan, Jason A Fries

Received: 1 April 2019 Accepted: 15 August 2019

Published online: 25 September 2019

imperfectly labeled training data. These weakly supervised methods can exceed the performance of traditional supervised machine learning while eliminating the cost of obtaining large quantities of hand-labeled data for training.<sup>18,19</sup> Unlike labeled clinical data, these heuristics, implemented as programmatic labeling functions, can be easily modified and shared across institutions, thus overcoming a major barrier to learning from large quantities of healthcare data while protecting patient privacy.

In this work, we demonstrate the use of such novel methods to identify real-world patient outcomes for the one of the most commonly implanted medical devices: the hip joint implant. More than one million joint replacements are performed every year in the United States and rates are increasing, including in patients under 65 years of age.<sup>20</sup> Primary joint replacements are expensive,<sup>21</sup> and implant failure incurs significant financial and personal cost for patients and the health care system. Recent recalls for metal-on-metal hip replacements<sup>22</sup> highlight the need for scalable, automated methods for implant surveillance.<sup>23</sup> While major adverse events such as deaths or revision surgery are typically reported, comprehensive device surveillance also needs to track outcomes that are reported less often, such as implant-related infection, loosening, and pain. Pain is a primary indication for undergoing a hip replacement, yet studies have found that more than 25% of individuals continue to experience chronic hip pain after replacement.<sup>24,25</sup> Pain is also an early indicator of the need for revision surgery,<sup>26</sup> and thus a key implant safety signal.

Motivated by the potential impact and scale of this surveillance problem, we applied weak supervision and deep learning methods using the Snorkel framework,<sup>19</sup> to identify reports of implant-related complications and pain from clinical notes. After validating these machine reading methods, we combined data from clinical notes with structured EHR data to characterize hip implant performance in the real world, demonstrating the utility of using such text-derived evidence. Working with the EHRs from a cohort of 6583 hip implant patients, we show that our methods substantially increase the number of identified reports of implant-related complications, that our findings complement previous post-approval studies, and that this approach could improve existing approaches to identify poorly performing implants.

## RESULTS

### Overview

We first describe the cohort analyzed in our study and report the performance of our implant extraction method and weakly supervised relation extraction methods in comparison to traditionally supervised models. We then compare the ability of our method to replicate the contents of the manually curated STJR joint implant registry. Lastly, we present our findings when combining the evidence produced by our extraction methods with structured data from EHRs to characterize hip implant performance in the real world, demonstrating the utility of such text processing methods.

### Cohort summary

We identified 6583 patients with records of hip replacement surgery. The hip replacement cohort was 55.6% female, with an average age at surgery of 63 years and average follow-up time after replacement of 5.3 years ( $\pm 2.1$  years). Of these 6,583 patients, 386 (5.8%) had a coded record of at least one revision surgery. For patients who had a revision surgery, the average age at primary replacement surgery was 57.9 years and average follow-up time was 10.5 years ( $\pm 3.0$  years).

**Table 1.** Precision, recall and F1 of our best performing entity and relation extraction methods

Category	Task Type	Precision	Recall	F1	Test set size
Implant manufacturer/model	Entity	96.3	98.5	97.4	500
Pain-anatomy	Relation	80.2	82.5	81.3	100
Implant-complication	Relation	82.6	61.1	70.2	633

### Performance of machine reading methods

The performance of our implant, complication, and pain extraction methods based on test sets of manually annotated records is summarized in Table 1. During adjudication of differences amongst annotators, there was very little disagreement on specific cases—the majority of differences were as a result of one annotator missing a relation or entity that another annotator had correctly identified. For implant-complications and pain-anatomy relations, Fleiss' kappa was 0.624 and 0.796, respectively, indicating substantial agreement amongst annotators.

Table 2 reports the performance of our relation extraction models compared against the baselines described in the methods. Using weakly labeled training data substantially improved performance in all settings. Weakly supervised LSTM models improved 9.2 and 24.6 F1 points over the soft majority vote baseline, with 17.7 and 29.4 point gains in recall. Directly training an LSTM model with just the 150 hand-labeled documents in the training set increased recall (+13.5 and +15.4 points) over the majority vote for both tasks, but lost 8.9 and 30.8 points in precision. Model performance for each subcategory of relation is enumerated in Supplementary Table 4. These results show that training using large amount of generated, imperfectly labeled data is beneficial in terms of recall, while preserving precision.

For subcategories of implant complications, F1 was improved by 16.6–486.2% for different categories except radiographic abnormalities, where a rule-based approach performed 12.6% better. In general, models trained on imperfectly labeled data provided substantial increases in recall, especially in cases such as particle disease, where rules alone had very low recall.

Figure 1 shows the impact of training set size on the weakly supervised LSTM's performance for identifying implant complications. With training set sizes greater than 10,000 weakly labeled documents, LSTM models outperformed rules across all possible classification thresholds, with a 51.3% improvement in area under the precision-recall curve (AUPRC) over soft majority vote.

The STJR snapshot covered 3714 patients. 2850 STJR patients underwent surgery during the time period of this study (1995–2014). 1,877 STJR patients overlapped with our hip replacement cohort (Fig. 2) and were used to evaluate model performance. For the most frequently implanted systems with  $\geq 1$  records present in the STJR (Zimmer Biomet Trilogy, DePuy Pinnacle, Depuy Duraloc, Zimmer Biomet Continuum, Zimmer Biomet RingLoc, Zimmer Biomet VerSys, Depuy AML and Zimmer Biomet M/L Taper; corresponding to 88% of patients in the STJR/cohort overlapping set), there was 72% agreement between STJR and our system output, 17% conflict and 11% missingness.

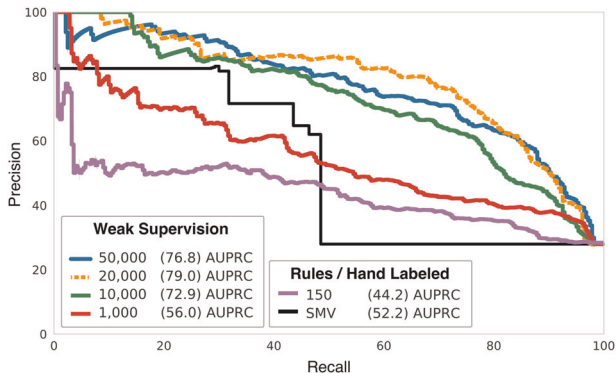
### Hip implant system performance in the real world

In the following sections we summarize our analysis of hip implant system performance using the evidence extracted by our machine reading methods. First we quantify additional evidence of revision identified by our methods, and the association between implant system and risk of revision. Next, we summarize our analysis of the association between implant systems and post-implant complications generally. Lastly, we describe our findings regarding the


**Table 2.** Performance in terms of precision (P), recall (R), and F1, with standard deviation (SD) for weakly supervised relation extraction compared with baselines

MODEL	Pain-Anatomy ( <i>n</i> = 236)			Implant-Complication ( <i>n</i> = 276)		
	P (SD)	R (SD)	F1 (SD)	P (SD)	R (SD)	F1 (SD)
Soft Majority Vote (SMV)	<b>81.4 (2.8)</b>	64.8 (3.0)	72.1 (2.3)	81.6 (3.6)	31.7 (2.7)	45.6 (3.1)
Fully Supervised (FS)	72.5 (2.9)	78.3 (2.6)	75.3 (2.1)	50.8 (3.1)	47.1 (3.1)	48.8 (2.7)
Weakly Supervised (WS)	80.2 (2.6)	<b>82.5 (2.4)</b>	<b>81.3 (1.9)</b>	<b>82.6 (2.6)</b>	<b>61.1 (2.9)</b>	<b>70.2 (2.3)</b>
Improvement over SMV	−1.5%	+27.3%	+12.8%	+1.2%	+92.7%	+53.9%

Bold highlights show highest value for a given metric



**Fig. 1** Precision Recall (AUPRC) curves for Implant-Complication classifier performance at different training set sizes. The soft majority vote (SMV) baseline consists of a majority vote of all labeling functions applied to test data. Note how SMV favors precision at substantial cost to recall. A classifier trained directly on the 150 hand-labeled training documents failed to produce a high-performance model. When training using the generated imperfectly labeled training data, at 10k–50k documents, the resulting supervised LSTMs are better than SMV at all threshold choices. The model trained on 20k documents performed best overall, with a 51.3% improvement over SMV in average precision (AUPRC)



**Fig. 2** Overlap between Stanford Total Joint Registry (STJR) and the hip replacement cohort extracted from the Stanford Medicine Research Data Repository (STaRR). Of the 6583 patients in the hip replacement cohort, 782 patients did not have an operative report (white box with dotted outline). For the 5801 patients with operative reports, 807 did not have a mention of any implant model in their report(s) (striped box). Of the remaining 4994 patients (light gray box), 1877 overlapped with the STJR. Of the 477 patients whose 500 operative reports were manually annotated (dark gray box), 185 overlapped with the STJR. 973 patient records were present only in the STJR (white box with dashed outline)

association between implant system, hip pain, and revision. Our first two analyses used Cox proportional hazards models to measure the association between implant system and revision/post-implant complications. Our third analysis used a negative


binomial model to investigate the association between implant system and frequency of hip pain mentions over time, and a t-test to investigate the association between hip pain and revision surgery.

We found that unstructured data captures additional evidence of revision. In the analyzed subset of 2704 patients with a single implant, only 78 had a coded record of revision surgery (as an ICD or CPT code in a billing record, for example). Considering only this evidence of revision, there was no significant association between implant system and revision-free survival (log-rank test  $p = 0.8$ ; see Supplementary Figure 2). When including evidence of revision extracted from clinical notes, we detected an additional 504 revision events (over six times as many events). In total, our methods identified 519 unique revision events (63 revision events that were detected in both coded data and text for a single patient were merged, and assigned the timestamp of whichever record occurred first). 63/78 (81%) of the coded revision events have corresponding evidence of the revision extracted from clinical notes by our system. The majority of coded records and text-derived revision events occur within 2 months of each other, suggesting good agreement between these two complementary sources of evidence when both do exist.

Analyzing the augmented revision data using a Cox proportional hazards model, different implant systems are associated with a significantly higher or lower risk of revision when controlling for age at the time of implant, race, gender, and Charlson Comorbidity Index (log-rank test  $p < 0.001$ ). Figure 3 summarizes the risk of revision for implant systems when including evidence from clinical notes alongside structured records of revision.

In addition to revision, the Cox proportional hazards analysis of post-hip replacement complications indicates that some implant systems are associated with a higher or lower risk of implant-related complications overall (log-rank test  $p < 0.001$ ) and of specific types (see Supplementary Figs. 3–7, for specific complication types). For example, the DePuy Pinnacle (acetabular component) + AML (femoral component) system has a significantly lower risk of overall complications in comparison to the Zimmer Biomet Trilogy (acetabular) + VerSys (femoral) system, when controlling for age at the time of implant, race, gender, and Charlson Comorbidity Index (see Fig. 4).

We also found that post replacement hip pain is associated with implant systems and revision. Our system extracted 938,222 positive pain-anatomy relation mentions from 471,985 clinical notes for 5562 hip replacement patients, with a mean of 168 mentions per patient, and 1 mention per note. Fitting a negative binomial model to the subset of this data corresponding to the 2704 patients with a single implant, we found that some implant systems are associated with higher frequency of hip pain mentions in the year following hip replacement surgery. Other covariates that also associated with higher frequency of post-implant hip pain are hip pain in the year prior to hip replacement,



**Fig. 3** Summary of Cox proportional hazards analysis of the risk of revision for each hip implant system. The table on the left lists the number of patients implanted with each system, the number of revision events observed for each, and the total person-years of data available. The forest plot displays the corresponding hazard ratio, with the hazard ratio (95% confidence interval) and p-value listed in the table to the right. Note that this figure only shows implant systems for which at least one revision event was detected

having the race category “other”, unknown ethnicity, and follow-up time (see Supplementary Table 1 for all model coefficients).

Table 3 lists implant systems with statistically significant negative binomial model coefficients ( $p \leq 0.05$ ), their incidence rate ratios (IRR) and 95% confidence intervals. Six systems, all manufactured by Depuy, had IRRs  $< 1$ , indicating that they are associated with lower rates of hip pain mentions relative to the Zimmer Biomet Trilogy + VerSys reference system when controlling for patient demographics, pain mentions in the prior year and CCI. Five systems (4 Zimmer Biomet, 1 Depuy) have IRRs  $> 1$ , indicating that they are associated with higher rates of hip pain mentions following implant.

For 1463 patients who had a BMI measurement within 100 days of hip replacement, we assessed whether hip pain mentions in the year following implant varies with BMI by including it in a negative binomial model alongside the above variables. We found no statistically significant association between BMI and pain in the year following hip replacement ( $p = 0.29$ ).


Mean post-hip replacement hip pain mention frequency is significantly higher in patients who underwent revision surgery, when controlling for length of post-implant follow-up time (8.94 vs. 3.23;  $t = 17.60$ ;  $p < 0.001$ ). This holds true when considering only the period between primary hip replacement and revision for those patients who had revision surgery (4.97 vs. 3.23;  $t = 5.14$ ;  $p < 0.001$ ).

## DISCUSSION

Post-market medical device surveillance in the United States currently relies on spontaneous reporting systems such as the FDA’s Manufacturer and User Facility Device Experience (MAUDE), and increasingly on device registries.<sup>36–39</sup> Neither source provides a complete or accurate profile of the performance of medical devices in the real world.<sup>23</sup> Submitting data to registries is voluntary and not all records capture complete details on primary procedure, surgical factors, complications, comorbidities, patient-reported outcomes, and radiographs,<sup>40</sup> so existing state and health system-level registries are not comprehensive. The shortcomings of spontaneous reporting systems, which are well known in pharmacovigilance research<sup>41,42</sup> (including lack of timeliness, bias in reporting, and low reporting rates), also plague device surveillance efforts.<sup>43</sup> EHRs, with their continuous capture of data from diverse patient populations and over long periods of time, offer a valuable source of evidence for medical device surveillance in the real world, and complement these existing resources. Indeed, the FDA’s new five year strategy<sup>44</sup> for its Sentinel post-market surveillance system prioritizes increased capture of data from EHRs. Our methods to extract evidence for device safety signal detection from EHRs directly support these efforts.

Our inference-based methods for identifying the implant manufacturer/model, pain and complications from the EHR were very accurate (F1 97.4; 81.3; 70.2) and substantially outperformed





**Fig. 4** Summary of Cox proportional hazards analysis of the risk of any complication for each hip implant system. The table on the left lists the number of patients implanted with each system, the number of complication events observed for each, and the total person-years of data available. The forest plot displays the corresponding hazard ratio, with the hazard ratio (95% confidence interval) and p-value listed in the table to the right. Note that this figure only shows implant systems for which one complication event was observed

**Table 3.** Negative binomial model-derived incidence rate ratios (IRRs), 95% confidence intervals (CI) and associated p-values for implant systems significantly associated with frequency of hip pain mentions in the year following hip replacement

Implant System	IRR (95% CI)	p-value
Depuy Duraloc + AML	0.033 (0.012–0.091)	<0.001
Depuy Duraloc + Corail	1.268 (1.020–1.577)	0.033
Depuy Duraloc + Summit	0.161 (0.040–0.652)	0.011
Depuy Pinnacle + AML	0.097 (0.057–0.163)	<0.001
Depuy Pinnacle + Corail	0.594 (0.475–0.742)	<0.001
Depuy Pinnacle + Endurance	0.032 (0.010–0.137)	<0.001
Depuy Pinnacle + Summit	0.369 (0.301–0.432)	<0.001
Zimmer Biomet Continuum + M/L Taper	2.061 (1.561–2.720)	<0.001
Zimmer Biomet Ranawat/Burstein + Taperloc	2.061 (1.174–3.621)	0.012
Zimmer Biomet Trilogy + M/L Taper	1.490 (1.234–1.799)	<0.001

pattern and rule-based methods. As a result, our approach improved coverage by identifying more patients who underwent hip replacement, and detecting subsequent complications in a more comprehensive manner. Moreover, these machine learning

methods can be deployed continuously, enabling near real-time automated surveillance. Comparing our method's output to a manually curated joint registry, we achieved majority agreement with existing registry records and increased coverage of hip replacement patients by 200%. Discrepancies between implant information extracted from operative reports and the corresponding registry record primarily resulted from variability in how implant system names were recorded; thus, one way our methods can augment existing registries is to standardize record capture. Our findings highlight the importance of using evidence derived from clinical notes in evaluating device performance. Our methods identified significantly more evidence of revision surgery (over 6X more events than from codes alone), augmenting what is available in structured EHR data. Our finding that text-derived evidence increases the observed rate of complications in comparison to coded data agrees with studies of the completeness and accuracy of ICD coding for a variety of diseases,<sup>45</sup> and for comorbidities and complications of total hip arthroplasty specifically.<sup>46,47</sup>

We found that a subset of implant systems were associated with higher or lower risk of implant-related complications in general, and of specific classes of complication. A recent meta-analysis of implant combinations<sup>48</sup> found no association between implant system and risk of revision, although other studies have mixed conclusions.<sup>49–54</sup> These studies relied on structured records of

revision and registry data, and thus are complemented by our analysis of real world evidence derived from EHRs. We also found that patients with structured records of revision surgery reported more hip pain in the post replacement, pre-revision period than patients who did not. This agrees with previous findings<sup>36,55</sup> that pain is an early warning sign of complications that result in revision surgery.

Our approach leverages data programming to generate training sets large enough to take advantage of deep learning<sup>32</sup> for relational inference (which is one of the most challenging problems in natural language processing<sup>27</sup>). The resulting models sacrificed small amounts of precision over a rule-based approach for significant gains in recall, demonstrating their ability to capture complex semantics and generalize beyond input heuristics. This approach requires only a small collection of hand-labeled data to validate end model performance. By focusing on creating labeling functions, instead of manually labeling training data, we achieved state-of-the-art performance with reusable code that, unlike labeled data, can be easily updated and shared across different healthcare systems. The generative model learned from data labeled using programmatic labeling functions may be applied to new datasets with zero-to-low hand-labeling effort required, and used to train high performing downstream discriminative models over data from different sites or sources. Prior work<sup>19,56</sup> demonstrates the utility of weak supervision in such variable and low-label settings, commonly encountered in medical applications.

There are several limitations to our work. We were not able to retrieve operative reports for all hip replacement patients and thus likely missed some implants. Patients had different numbers of clinical notes, at varying time points, which may not have captured their full experience of pain or complications following surgery. It is also possible that the modes of note entry available in different EHR systems may encourage more detailed reporting of complications in different health systems, which would affect the information our system extracts. Our approach canonicalized implants to the level of manufacturer/model names, which only identifies broad implant families rather than the specific details provided by serial and lot numbers. However, as the use of unique device identifiers grows in popularity, our approach can incorporate this information to further differentiate the performance of implants with different design features. Our study was necessarily restricted to the implant systems used by Stanford Health Care surgeons, and thus our findings are limited to those systems. While we controlled for patient age, gender, race, ethnicity, hip pain mentions prior to surgery, and CCI in our statistical analyses, there are additional patient- and practice-specific features that may confound the complication-free survival of individual implants and rates of post-implant hip pain. These include case complexity, the orthopaedic surgeon who performed the hip replacement, their preference for specific implant systems, and their case mix. CCI is an indicator of case complexity, but joint arthroplasty-specific factors, such as indication for surgery, also contribute to patients' overall complexity. Therefore, we focus on developing scalable methods that may be deployed at other sites to develop a national-level device registry populated from EHRs (we make our code publicly available), and restrict the analysis of implant complications to demonstrating an improvement in inference over using only registry collected data. With the resulting larger dataset it may be feasible to use advanced methods such as propensity score modeling to select matched cohorts for analysis that account for patient-level variation and their association with likelihood of receiving a specific implant system, but this is outside the scope of the current work and not currently possible with the single-site data we have access to.

Lastly, our analysis was restricted to patients who underwent a single hip replacement, to avoid attributing a complication or pain event to the incorrect implant, in the case of patients who

underwent multiple primary implant procedures. Given these limitations, the challenge of obtaining perfect test data (as evidenced by the observed intra-annotator variability, and noise in ground truth labels from the STJR), and the loss of precision introduced by weak supervision (in exchange for significant gains in recall), we would emphasize our findings are exploratory, and would be strengthened through additional studies.

In conclusion, we demonstrate the feasibility of a scalable, accurate, and efficient approach for medical device surveillance using EHRs. We have shown that implant manufacturer and model, implant-related complications, as well as mentions of post-implant pain can be reliably identified from clinical notes in the EHR. Leveraging recent advances in machine reading and deep learning, our methods require orders of magnitude less labeled training data and obtain state-of-the-art performance. Our findings that implant systems vary in their revision-free survival and that patients who had revision surgery had more mentions of hip pain after their primary hip replacement agree with multiple prior studies. Associations between implant systems, complications, and hip pain mentions as found in our single-site study demonstrate the feasibility of using EHRs for device surveillance, but do not establish causality. The ability to quantify pain and complication rates automatically over a large patient population offers an advantage over surveillance systems that rely on individual reports from patients or surgeons. We believe that the algorithms described here can be readily scaled, and we make them freely available for use in analyzing electronic health records nationally.

## METHODS

### Overview

We developed machine reading methods to analyze clinical notes and identify the implanted device used for a patient's hip replacement, as well as identify mentions of implant-related complications and patient-reported pain. We evaluated these methods' ability to (1) accurately identify implant-related events with a minimum of hand-labeled training data; and (2) map identified implants to unique identifiers—a process called canonicalization—to automatically replicate the contents of a manually curated joint implant registry.

We then combined the data produced by our machine learning methods with structured data from EHRs to: (1) compare complication-free survivorship of implant systems and (2) derive insights about associations between complications, pain, future revision surgery and choice of implant system.

In the following sections, we first describe how we identified the patient cohort analyzed in our study. We then describe in detail our novel weakly supervised machine reading methods. Lastly, we describe our analysis of the performance of implant systems based on real-world evidence produced by these methods, using Cox proportional hazards models, negative binomial models, and the t-test.

Our study was approved by the Stanford University Institutional Review Board with waiver of informed consent, and carried out in accordance with HIPAA guidelines to protect patient data.

### Cohort construction

We queried EHRs of roughly 1.7 million adult patients treated at Stanford Health Care between 1995 and 2014 to identify patients who underwent primary hip replacement and/or revision surgery. To find records of primary hip replacement surgery, we used ICD9 procedure code 81.51 (total hip replacement) and CPT codes 27130 (total hip arthroplasty) and 27132 (conversion of previous hip surgery to total hip arthroplasty). To find structured records of hip replacement revision surgery, we used ICD9 procedure codes 81.53 (Revision of hip replacement, not otherwise specified), 00.70, 00.71, 00.72, and 00.73 (Revision of hip replacement, components specified), and CPT codes 27134 (Revision of total hip arthroplasty; both components, with or without autograft or allograft), 27137 (Revision of total hip arthroplasty; acetabular component only, with or without autograft or allograft) and 27138 (Revision of total hip arthroplasty; femoral component only, with or without allograft).

We identified 6583 patients with records of hip replacement surgery, of which 386 (5.8%) had a coded record of at least one revision after the primary surgery. For all patients in the resulting cohort, we retrieved the entirety of their structured record (procedure and diagnosis codes, medication records, vitals etc.), their hip replacement operative reports, and all clinical notes. We were able to retrieve operative reports for 5801 (88%) hip replacement patients.

### Machine reading

In supervised machine learning, experts employ a wide range of domain knowledge to label ground truth data. Data programming<sup>18,19</sup> is a method of capturing this labelling process using a collection of imperfect heuristics or labeling functions which are used to build large training sets. Labeling functions encode domain expertise and other sources of indirect information, or weak supervision (e.g., knowledge bases, ontologies), to vote on a data item's possible label. These functions may overlap, conflict, and have unknown accuracies. Data programming unifies these noisy label sources, using a generative model to estimate and correct for the unobserved accuracy of each labeling source and assign a single probabilistic label to each unlabeled input sample. This step is unsupervised and requires no ground truth data. The generative model is used to programmatically label a large training set, to learn a discriminative model such as a neural network. By training a deep learning model, we transform rules into learned feature representations, which allow us to generalize beyond the original labeling heuristics, resulting in improved classification performance.

The probabilistic labels generated via data programming results in a large training set, which is referred to as “weakly labeled”. The discriminative model gets as input the original text sentences in the training set and the probabilistic labels generated by the data programming step. Because the label provided in such a supervised learning setup is “weak”, the process is also called “weak supervision” and the resulting discriminative model referred to as a weakly supervised model. We describe each step of the process (Fig. 5) in the following paragraphs.

We defined 3 entity/event types: (1) implant system entities identified by a manufacturer and/or model name, e.g., “Zimmer VerSys”; (2) implant-related complications, e.g., “infected left hip prosthetic”; and (3) patient-reported pain at a specific anatomical location, e.g., “left hip tenderness”. The latter two, implant complications and patient-reported pain, involve multiple concepts and were formulated as relational inference tasks<sup>27</sup> where a classifier predicts links between two or more entities or concepts found anywhere in a sentence, i.e.,  $R = r(c_1, \dots, c_N)$ , where  $r$  is a relation type and  $c_i$  is a concept mention. This formulation allowed our classifier learn complex semantic relationships between sentence entities, using statistical inference to capture a wide range of meaning in clinical notes. This approach enabled identifying nuanced, granular information, such as linking pain and complication events to specific anatomical locations and implant subcomponents (e.g., “left acetabular cup demonstrates extreme

liner wear”). Therefore, complication outcomes were represented as entity pairs of (complication, implant) and pain outcomes as pairs of (pain, anatomy).

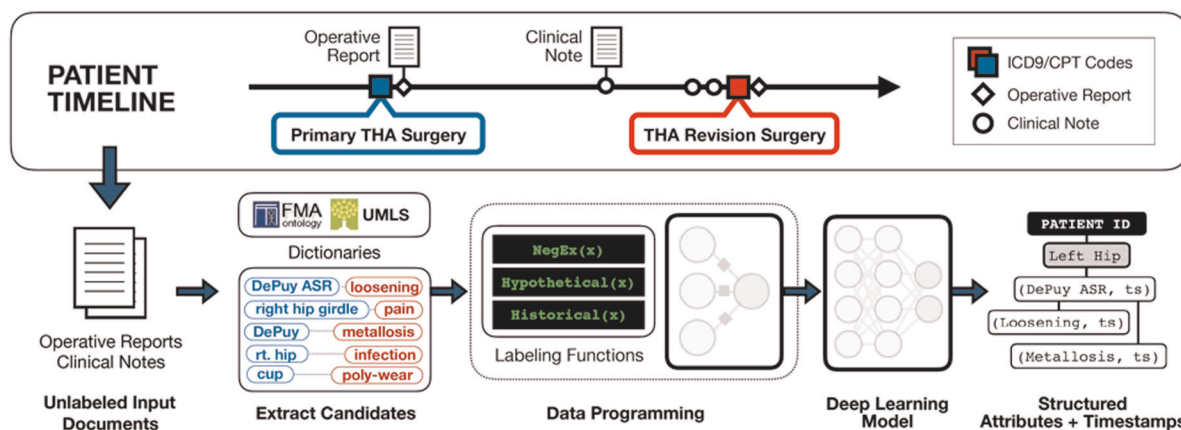
We considered all present positive mentions of an event—i.e., in which the event was contemporaneous with a note's creation timestamp—as positive examples. Historical, negated and hypothetical mentions were labeled as ‘negative’. Implant complications were further broken down into 6 disjoint subcategories: revision, component wear, mechanical failure, particle disease, radiographic abnormality, and infection (see Supplementary Material for examples). These categories correspond to the removal of specific hardware components as well as common indicators of device failure.

All clinical notes were split into sentences and tokenized using the spaCy<sup>28</sup> framework. Implant, complication, anatomy, and pain entities were identified using a rule-based tagger built using a combination of biomedical lexical resources, e.g., the Unified Medical Language System (UMLS)<sup>29</sup> and manually curated dictionaries (see supplement for details). Each sentence was augmented with markup to capture document-level information such as note section headers (e.g., “Patient History”) and all unambiguous date mentions (e.g., “1/1/2000”) were normalized to relative time delta bins (e.g., “0–1 days”, “+ 5 years”) based on document creation time. Such markup provides document-level information to incorporate into features learned by the final classification model.

Candidate relations were defined as the Cartesian product of all entity pairs defining a relation (i.e., pain/anatomy and implant/complication) for all entities found within a sentence. Candidate events and relations were generated for all sentences in all clinical notes.

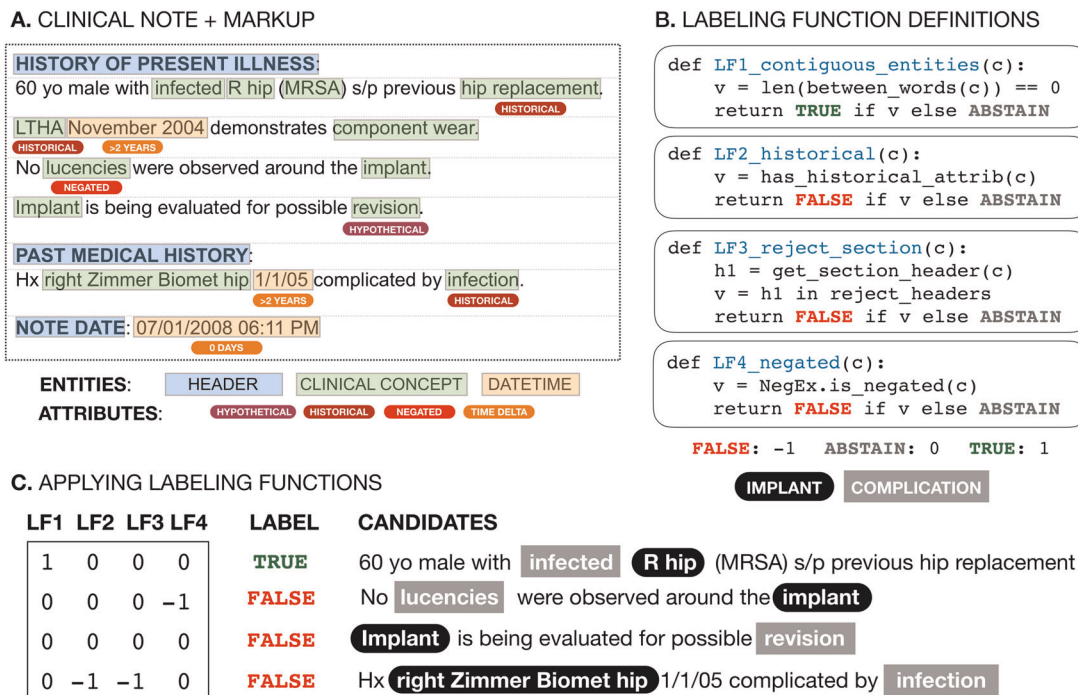
Labeling functions leverage existing natural language tools to make use of semantic abstractions, or primitives, which make it easier for domain experts to express labeling heuristics. For example, primitives include typed named entities such as bacterium or procedures, document section structure, as well as negation, hypothetical, and historical entity modifiers provided by well-established clinical natural language processing methods such as NegEx/ConText.<sup>30,31</sup> These primitives provide a generalized mechanism for defining labeling functions that express intuitive domain knowledge, e.g., “Implant complications found in the ‘Past Medical History’ section of a note are not present positive events” (Fig. 6). Primitives can be imperfect and do not need to be directly provided to the final discriminative model as features. Critically, the end deep learning discriminative model automatically learns to encode primitive information using only the original text and the training signal provided by our imperfect labels.

We developed labeling functions to identify implant-related complications and pain and its anatomical location. In total, 50 labeling functions were written, 17 shared across both tasks, with 7 task-specific functions for pain-anatomy and 25 for implant-complications. Labeling functions were written by inspecting unlabeled sentences and developing heuristics to assign individual labels (Supplementary Fig. 2). This process was iterative,



**Fig. 5** Overview of our machine reading pipeline. Top: Each patient's EHR is processed to extract the date of primary hip replacement surgery, any coded record of revision surgeries, and all clinical and operative notes. Bottom: From the patient's coded data and primary hip replacement operative report, we tagged all mentions of implants, complications, and anatomical locations. We defined pairs of relation candidates from these sets of entities, and labeled them using data programming via the Snorkel framework. These labeled data were then used to train a deep learning model. When applied to unseen data, the final model's final output consists of timestamped, structured attribute data for implant systems, implant-related complications, and mentions of pain





**Fig. 6** Labeling function examples. Clinical notes (top left) are preprocessed to generate document markup, tagging entities with clinical concepts (blue), parent section headers (green), dates/times (orange), and historical, hypothetical, negated, and time-delta attributes (shades of red). Labeling functions (top right) use this markup to represent domain insights as simple Python functions, e.g., a function to label mentions found in “Past Medical History” as FALSE because they are likely to be historical mentions rather than a current condition. These labeling functions vote (FALSE, ABSTAIN, TRUE) on candidate relationships (bottom, with true labels indicated in the LABEL column) consisting of implant (black rounded box) - complication (grey square box) pairs to generate a vector of noisy votes for each candidate relationship. The data programming step described in Fig. 1 uses this labeling function voting matrix to learn a generative model which assigns a single probabilistic label to each unlabeled candidate relationship

using development set data (described in the next section) to refine labeling function design.

We used two sources of manually annotated data to evaluate our methods: (1) the Stanford Total Joint Registry (STJR), a curated database of joint replacement patients maintained by Stanford Health Care orthopedic surgeons which, as of March 2017, contained records for 3714 patients and (2) clinical notes from the Stanford Health Care EHR system, manually annotated by co-authors and medical doctors expressly for the purpose of evaluating our models’ performance. The clinical notes manually annotated by our team consisted of three sets:

- 500 operative reports annotated to identify all implant mentions.
- 802 clinical notes annotated to identify all complication-implant mentions.
- 500 clinical notes annotated to identify all pain-anatomy relation mentions.

Set A was randomly sampled from the entire hip replacement cohort’s operative report corpus, using the manufacturer search query “Zimmer OR DePuy”. This query captured >85% of all implant components in the STJR and accounted for >90% of all implant mentions in operative reports. Sets B and C were randomly sampled with uniform probability from the entire hip replacement cohort’s clinical note corpus. Ground truth labels for implant model mentions (set A) were generated by one annotator (author AC). Ground truth labels for complication-implant mentions (set B) were generated by five annotators (authors AC, JF, and NHS, and two medical doctors). Ground truth labels for set C were generated by three annotators (authors AC and JF, and one medical doctor). We used adjudication to resolve differences between annotators.

For implant manufacturer/model entity extraction (e.g., “DePuy Pinna- cle”, “Zimmer Longevity”), we used a rule-based tagger, as we found using dictionary-based string matching was sufficiently unambiguous to achieve high precision and recall (see supplemental methods for detail on the construction of the implant dictionary used).

We used a Bidirectional Long Short-Term Memory (LSTM)<sup>32</sup> neural network with attention<sup>33</sup> as the discriminative model for relational

inference tasks. All LSTMs used word embeddings pre-trained using FastText<sup>34</sup> on 8.1B tokens (651 M sentences) of clinical text. Each training instance consisted of an entity pair, its source sentence, and all corresponding text markup. All models were implemented using the Snorkel framework.

We evaluated our implant extraction system using the 500 manually annotated operative reports (set A, above). We evaluated our weakly supervised pain and complication LSTMs against two baselines: (1) a traditionally supervised LSTM trained using 150 hand-labeled training documents; (2) the soft majority vote of all labeling functions for a target task. To quantify the effect of increasing training set size, we evaluated the weakly supervised LSTM neural networks trained using 150 to 50,000 weakly labeled documents. We used training, development, and test set splits of 150/19/633 from set B for implant-complication relations and 150/250/100 documents from set C for pain-anatomy relations. All models and labeling functions were tuned on the development split and results reported are for the test split. All models were scored using precision, recall, and F1-score.

We then compared the structured output of our implant extraction method to the STJR. Given a patient, an implant component, and a timestamped hip replacement surgery, agreement was defined as identical entries for the component in both the STJR and our system’s output; conflict was where the implant component(s) were recorded for a hip replacement surgery in both STJR and our data but did not match; missingness was where a record for a hip replacement surgery was absent in either STJR or our data. We canonicalized all implant models to the level of manufacturer/model, i.e., the same resolution used by the STJR.

### Characterizing hip implant performance in the real world

We analyzed complication-free survival of implant systems using Cox proportional hazards models, controlling for age at the time of hip replacement, gender, race, ethnicity and Charlson Comorbidity Index (CCI). CCI was categorized as none (CCI = 0), low (CCI = 1), moderate (CCI = 2), or high (CCI ≥ 3). For patients with multiple complications, we calculated



survival time from primary hip replacement to first complication. We performed this analysis for a composite outcome of 'any complication', and also for each class of complication extracted by our text processing system.

To quantify association between implant systems and hip pain, we fit a negative binomial model to the frequency of hip pain mentions in the year post-hip replacement. The model included the following covariates: implant system, age at the time of hip replacement, gender, race, ethnicity, frequency of hip pain mentions in the year prior to surgery (to account for baseline levels of hip pain), and follow-up time post-THA (to a maximum of one year). For the subset of patients with body mass index (BMI) data available, we also fit models that included BMI as a covariate.

We specified implant system as an interaction term between femoral and acetabular components, grouping infrequently occurring interactions into a category "other". The frequency cutoff for inclusion in the "other" category was chosen in a data-driven manner using Akaike information criterion (AIC) to find the cutoff resulting in best model fit. We specified the most frequent system (Zimmer Trilogy acetabular + VerSys femoral component) as the reference interaction term.

To assess the association between revision and hip pain, we tested the hypothesis that hip-specific pain mentions were more frequent in patients with a coded record of revision compared to those who without. We used the two-sided t-test, controlling for post-implant follow-up time.

For all implant system analyses, we restricted the analysis to the 2704 patients with a single hip replacement, a single femoral component and a single acetabular component.

All statistical analyses were performed using R 3.4.0.

### Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this article.

### DATA AVAILABILITY

The data that support the findings of this study are available on request from the corresponding author A.C. The data are not publicly available because they contain information that could compromise patient privacy.

### CODE AVAILABILITY

Code, dictionaries, and labeling function resources required to extract implant details, complication-implant and pain-anatomy relations are publicly available at <https://github.com/som-shahlab/ehr-rwe>. Included are Jupyter Notebooks that run the complication-implant and pain-anatomy inference code end-to-end on MIMIC-III<sup>35</sup> notes.

### ACKNOWLEDGEMENTS

The authors acknowledge the contributions of Dr. Saurabh Gombard MD, Dr. Husham Sharifi MD, and Dr. James Tooley MD, who annotated clinical notes for mentions of implant-related complications, pain and anatomical location to enable evaluation of our text processing methods, and Dr. Kenneth Jung who provided guidance in carrying out statistical analyses. This work was supported in part by the National Institutes of Health National Library of Medicine through Grant R01LM011369, and the Mobilize Center, a National Institutes of Health Big Data to Knowledge (BD2K) Center of Excellence supported through Grant U54EB020405.

### AUTHOR CONTRIBUTIONS

A.C. and J.A.F. are co-first authors. A.C. and N.H.S. conceived the project. A.C. and J.A.F. developed and evaluated the machine reading and text processing methods used. J.I.H. provided access to joint registry data for methods evaluation. A.C. and N.H.S. planned statistical analyses, and A.C. performed all statistical analyses. C.R. provided computer science and informatics expertise. J.I.H. and N.G. provided clinical expertise. S.D. provided mentorship to J.A.F. A.C. wrote the first draft of the paper and all authors revised it. All authors have approved this version for publication.

### ADDITIONAL INFORMATION

**Supplementary Information** accompanies the paper on the *npj Digital Medicine* website (<https://doi.org/10.1038/s41746-019-0168-z>).

**Competing interests:** The authors declare no competing interests.

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

### REFERENCES

- Ibrahim, A. M. & Dimick, J. B. Monitoring medical devices: missed warning signs within existing data. *JAMA* **318**, 327–328 (2017).
- Resnic, F. S. & Normand, S.-L. T. Postmarketing surveillance of medical devices—filling in the gaps. *N. Engl. J. Med.* **366**, 875–877 (2012).
- Rising, J. & Moscovitch, B. The Food and Drug Administration's unique device identification system: better postmarket data on the safety and effectiveness of medical devices. *JAMA Intern. Med.* **174**, 1719–1720 (2014).
- Drozda, J. P. Jr et al. Constructing the informatics and information technology foundations of a medical device evaluation system: a report from the FDA unique device identifier demonstration. *J. Am. Med. Inform. Assoc.* **25**, 111–120 (2018).
- Woodman, B. S. et al. The Implant Files: a global investigation into medical devices - ICJ. ICJ (2018). <https://www.icj.org/investigations/implant-files/>. (2018).
- Blumenthal, D. & Tavenner, M. The 'meaningful use' regulation for electronic health records. *N. Engl. J. Med.* **363**, 501–504 (2010).
- Safran, C. et al. Toward a national framework for the secondary use of health data: an American Medical Informatics Association White Paper. *J. Am. Med. Inform. Assoc.* **14**, 1–9 (2007).
- Friedman, C. P., Wong, A. K. & Blumenthal, D. Achieving a nationwide learning health system. *Sci. Transl. Med.* **2**, 57cm29 (2010).
- Harpaz, R. et al. Combining signals from spontaneous reports and electronic health records for detection of adverse drug reactions. *J. Am. Med. Inform. Assoc.* **20**, 413–419 (2013).
- LePendur, P. et al. Pharmacovigilance using clinical notes. *Clin. Pharmacol. Ther.* **93**, 547–555 (2013).
- Capurro, D., Yetisgen, M., Eaton, E., Black, R. & Tarczy-Hornoch, P. Availability of structured and unstructured clinical data for comparative effectiveness research and quality improvement: a multi-site assessment. *eGEMS* **2**, 11 (2014).
- Lv, X., Guan, Y., Yang, J. & Wu, J. Clinical relation extraction with deep learning. *IJHIT* **9**, 237–248 (2016).
- Wu, Y., Jiang, M., Lei, J. & Xu, H. Named entity recognition in chinese clinical text using deep neural network. *Stud. Health Technol. Inform.* **216**, 624–628 (2015).
- Fries, J. A. Brundlefly at SemEval-2016 Task 12: Recurrent Neural Networks vs. Joint Inference for Clinical Temporal Information Extraction. Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016. 1274–1279 (2016).
- Jagannatha, A. N. & Yu, H. Bidirectional RNN for Medical Event Detection in Electronic Health Records. in **2016**, 473–482 (2016).
- Ravi, D. et al. Deep Learning for Health Informatics. *IEEE J. Biomed. Health Inf.* **21**, 4–21 (2017).
- Esteve, A. et al. A guide to deep learning in healthcare. *Nat. Med.* **25**, 24–29 (2019).
- Ratner, A., De Sa, C., Wu, S., Selsam, D. & Christopher, R. Data Programming: Creating Large Training Sets, Quickly. arXiv [stat.ML] (2016).
- Ratner, A. et al. Snorkel: rapid training data creation with weak supervision. *Proc. VLDB Endow.* **11**, 269–282 (2017).
- Kurtz, S. M. et al. Future young patient demand for primary and revision joint replacement: national projections from 2010 to 2030. *Clin. Orthop. Relat. Res.* **467**, 2606–2612 (2009).
- Lam, V., Teutsch, S. & Fielding, J. Hip and knee replacements: a neglected potential savings opportunity. *JAMA* **319**, 977–978 (2018).
- Cohen, D. Out of joint: The story of the ASR. *BMJ* **342**, d2905 (2011).
- Resnic, F. S. & Matheny, M. E. Medical devices in the real world. *N. Engl. J. Med.* **378**, 595–597 (2018).
- Nikolajsen, L., Brandsborg, B., Lucht, U., Jensen, T. S. & Kehlet, H. Chronic pain following total hip arthroplasty: a nationwide questionnaire study. *Acta Anaesthesiol. Scand.* **50**, 495–500 (2006).
- Erlenwein, J. et al. Clinical relevance of persistent postoperative pain after total hip replacement—a prospective observational cohort study. *J. Pain. Res.* **10**, 2183–2193 (2017).
- Eneqvist, T., Nemes, S., Bülow, E., Mohaddes, M. & Rolfson, O. Can patient-reported outcomes predict re-operations after total hip replacement? *Int. Orthop.* <https://doi.org/10.1007/s00264-017-3711-z> (2018).
- Mintz, M., Bills, S., Snow, R. & Jurafsky, D. Distant supervision for relation extraction without labeled data. *Association for Computational Linguistics.* **2**, 1003–1011 (2009).
- Honnibal, M. & Montani, I. spaCy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing [Computer software]. (2019). Retrieved from <https://spacy.io/>
- Lindberg, D. A., Humphreys, B. L. & McCray, A. T. The unified medical language system. *Methods Inf. Med.* **32**, 281–291 (1993).

30. Chapman, W. W., Bridewell, W., Hanbury, P., Cooper, G. F. & Buchanan, B. G. A simple algorithm for identifying negated findings and diseases in discharge summaries. *J. Biomed. Inform.* **34**, 301–310 (2001).
31. Chapman, W. W., Chu, D. & Dowling, J. N. ConText: An Algorithm for Identifying Contextual Features from Clinical Text. in Proceedings of the Workshop on BioNLP 2007: Biological, Translational, and Clinical Language Processing 81–88 (Association for Computational Linguistics 2007).
32. Zhou, P. et al. Attention-based bidirectional long short-term memory networks for relation classification. in Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers) **2**, 207–212 (2016).
33. Xu, K. et al. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. in International Conference on Machine Learning 2048–2057 (2015).
34. Joulin, A., Grave, E., Bojanowski, P. & Mikolov, T. Bag of Tricks for Efficient Text Classification. arXiv [cs.CL] (2016).
35. Johnson, A. E. W. et al. MIMIC-III, a freely accessible critical care database. *Sci. Data* **3**, 160035 (2016).
36. Franklin, P. D., Allison, J. J. & Ayers, D. C. Beyond joint implant registries: a patient-centered research consortium for comparative effectiveness in total joint replacement. *JAMA* **308**, 1217–1218 (2012).
37. Paxton, E. W., Inacio, M. C. & Kiley, M.-L. The Kaiser Permanente implant registries: effect on patient safety, quality improvement, cost effectiveness, and research opportunities. *Perm. J.* **16**, 36–44 (2012).
38. Hughes, R. E., Hallstrom, B. R., Cowen, M. E. & Igrisan, R. M. Michigan Arthroplasty Registry Collaborative Quality Initiative (MARCQI) as a model for regional registries in the United States. *Orthop. Res. Rev.* **2015**, 47–56 (2015).
39. Hughes, R. E., Batra, A. & Hallstrom, B. R. Arthroplasty registries around the world: valuable sources of hip implant revision risk data. *Curr. Rev. Musculoskelet. Med.* **10**, 240–252 (2017).
40. Hansen, V. J. et al. Registries collecting level-I through IV Data: institutional and multicenter use: AAOS exhibit selection. *J. Bone Jt. Surg. Am.* **96**, e160 (2014).
41. Stephenson, W. P. & Hauben, M. Data mining for signals in spontaneous reporting databases: proceed with caution. *Pharmacoepidemiol. Drug Saf.* **16**, 359–365 (2007).
42. Bate, A. & Evans, S. J. W. Quantitative signal detection using spontaneous ADR reporting. *Pharmacoepidemiol. Drug Saf.* **18**, 427–436 (2009).
43. Bates, J. et al. Quantifying the utilization of medical devices necessary to detect postmarket safety differences: A case study of implantable cardioverter defibrillators. *Pharmacoepidemiol. Drug Saf.* <https://doi.org/10.1002/pds.4565> (2018).
44. U.S. Food and Drug Administration. Sentinel System – Five-Year Strategy 2019–2023. (2019).
45. Wei, W.-Q. et al. Combining billing codes, clinical notes, and medications from electronic health records provides superior phenotyping performance. *J. Am. Med. Inform. Assoc.* **23**, e20–7 (2016).
46. Bozic, K. J. et al. Is administratively coded comorbidity and complication data in total joint arthroplasty valid? *Clin. Orthop. Relat. Res.* **471**, 201–205 (2013).
47. Mears, S. C. et al. Coding of diagnoses, comorbidities, and complications of total hip arthroplasty. *Clin. Orthop. Relat. Res.* **402**, 164–170 (2002).
48. López-López, J. A. et al. Choice of implant combinations in total hip replacement: systematic review and network meta-analysis. *BMJ* **359**, j4651 (2017).
49. Abdulkarim, A., Ellanti, P., Motterlini, N., Fahey, T. & O'Byrne, J. M. Cemented versus uncemented fixation in total hip replacement: a systematic review and meta-analysis of randomized controlled trials. *Orthop. Rev.* **5**, e8 (2013).
50. Paxton, E. et al. Risk of revision following total hip arthroplasty: metal-on-conventional polyethylene compared with metal-on-highly cross-linked polyethylene bearing surfaces: international results from six registries. *J. Bone Jt. Surg. Am.* **96**, 19–24 (2014).
51. Mäkelä, K. T. et al. Failure rate of cemented and uncemented total hip replacements: register study of combined Nordic database of four nations. *BMJ* **348**, f7592 (2014).
52. Nieuwenhuijse, M. J., Nelissen, R. G. H. H., Schoones, J. W. & Sedrakyan, A. Appraisal of evidence base for introduction of new implants in hip and knee replacement: a systematic review of five widely used device technologies. *BMJ* **349**, g5133 (2014).
53. Junnila, M. et al. Implant survival of the most common cemented total hip devices from the Nordic Arthroplasty Register Association database. *Acta Orthop.* **87**, 546–553 (2016).
54. Moskal, J. T., Capps, S. G. & Scanelli, J. A. Still no single gold standard for using cementless femoral stems routinely in total hip arthroplasty. *Arthroplast Today* **2**, 211–218 (2016).
55. Sueyoshi, T. et al. Clinical predictors for possible failure after total hip arthroplasty. *Hip Int.* **26**, 531–536 (2016).
56. Ratner, A., Hancock, B., Dunnmon, J., Sala, F., Pandey, S. & Ré, C. Training complex models with multi-task weak supervision. Proceedings of the AAAI Conference on Artificial Intelligence. **33**, (2019).



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019