

Using Work System Theory, Facets of Work, and Dimensions of Smartness to Characterize Applications and Impacts of AI

Steven Alter

University of San Francisco, San Francisco, CA 94117, USA

alter@usfca.edu

Abstract. This paper presents an approach for describing and characterizing algorithms that are discussed as though they embody artificial intelligence. After identifying key assumptions related to algorithms and summarizing work system theory (WST), this paper uses a hypothetical example to introduce aspects of WST and two additional ideas, facets of work and dimensions of smartness in devices and systems. Next, it applies those ideas to aspects of five AI-related examples presented by entrepreneurs and researchers at an MIT AI conference in July 2020. Those examples were selected because they illustrated many AI-related issues. This paper's contribution is a new approach for characterizing real world applications and impacts of almost any system that uses algorithms or is associated with artificial intelligence.

Keywords: Artificial Intelligence, Algorithm, Work System Theory, Facets of Work, Dimensions of Smartness

An abbreviated 12-page version of this paper appears in the Proceedings of the IFIP WG 8.6 Working Conference, Dec. 17-19, 2020, Volume 617 of the IFIP Advances in Information and Communication Technology series

1 Moving Beyond the Multiple Meanings of AI

Many discussions of AI in academia and in public venues revolve around vague definitions, cherry-picked examples, and a *mélange* of diverse opinions and observations from pundits and researchers whose comments are often taken out of context. Examples often fall into categories that are only tangentially related to each other: intelligent machines, neural networks, machine learning, expert systems, smart systems, cognitive computing, speech recognition, pattern recognition, image recognition, natural language processing, statistical algorithms, automated decision-making, and so on. Beyond various historical and technical commonalities, it is difficult to argue that those topics are all instances of a coherent and well-defined phenomenon rather than just a bunch of topics that are often thrown together under the umbrella term AI. The lack of clarity about what AI means makes it difficult to discuss whether the benefits, risks, and ethics of using AI differ in any significant way from the benefits, risks, and ethics of automation or computerization in general.

Goal and Assumptions. This paper explains how to use work system theory, facets of work, and dimensions of smartness for characterizing applications and impacts of AI. Those ideas will be explained through an illustrative example. This paper assumes that understanding affordances, benefits, and risks of AI applications in sociotechnical contexts can be viewed as a special case of concerns about algorithms, which was the central topic of the IFIP 8.2 working conference “Living with Monsters? Social Implications of Algorithmic Phenomena, Hybrid Agency, and the Performativity of Technology” [1]. Also, instead of speaking about algorithms in general, we look at algorithms (and hence AI) in the context of sociotechnical or totally automated work systems in which they are used.

Organization. First this paper defines key assumptions related to algorithms and then summarizes work system theory (WST). Given this paper’s length limitations, a hypothetical example of a hiring work system that uses several types of algorithms is used to introduce aspects of WST and two additional ideas, facets of work and dimensions of smartness in devices and systems. Next, this paper applies those in discussions of aspects of AI-related examples presented by entrepreneurs and researchers at an MIT AI conference in Silicon Valley in July 2020. This paper’s contribution is a new approach for characterizing applications and impacts of almost any system that uses algorithms, including systems whose algorithms are associated with artificial intelligence, big data, block chain, Internet of things, social media, and other current areas of interest associated with emerging technologies.

2 Assumptions Related to Algorithms

The examples in Table 1 use algorithms that may or may not be associated with AI. Some of those algorithms might be simple decision rules such as allowing no more than 40% of applicants to be classified in category X. Even a simple algorithm like that one can have other important and far reaching effects such as favoring one group of people over other groups, as when category X is acceptance into college. Examples

in Table 1 illustrate the difficulty of generalizing about benefits, risks, and ethics of AI without specifying the area of application.

Table 1. Potential application situations for algorithms that might or might not use AI

<ul style="list-style-type: none"> • using facial images to identify people • converting spoken words into equivalent text • deciding which applicants should be hired or accepted by a university • deciding whether to alert medical staff about a change in a patient’s condition • deciding whether a person has permission to drive a car • deciding which is the best target for a missile • deciding a person’s salary or bonus • deciding whether an autonomous (self-driving) vehicle needs to stop or swerve • controlling the aerodynamics of a rocket 	<ul style="list-style-type: none"> • deciding whether to turn off a machine likely to have a mechanical failure soon • deciding where police should be deployed over the next eight hours • selecting defective items that are being moved on a conveyor belt • combining multiple items in a customer order to minimize mailing cost • translating a text between languages • determining the best route for driving from the current location to a destination • finding the laws that are most relevant to a specific lawsuit
--	---

3 Assumptions Related to Algorithms

A series of assumptions are a starting point for visualizing and evaluating algorithms.

Algorithms as Specifications for Transforming Inputs into Outputs. An algorithm specifies exactly how to human and/or nonhuman actors can convert specific inputs into specific outputs. Algorithms are abstractions that cannot do anything on their own. Human and/or nonhuman actors perform the transformations.

Goals, Constraints, Other Parameters. Algorithms pursue goals, operate within constraints, and may be guided by other situation-specific parameters or inputs.

Omissions. Most algorithms that are not derived directly from mathematics have omissions, i.e., potentially important topics or issues that the algorithm ignores.

Biases. Most algorithms not derived directly from mathematics or theory bring purposeful or accidental biases. Those biases may come from the viewpoints of an algorithm’s creators or may be unintended results of omissions, biased training data, other shortcomings of the algorithm, or unanticipated interactions related to factors in the environment surrounding a work system in which the algorithm operates.

Areas of Greater and Lesser Acuity: Algorithms apply to specific domains, i.e., defined sets of things or conditions. Often they have areas of maximum relevance and acuity and other areas of limited relevance and acuity. Applying an algorithm near or just beyond the boundaries of its domain of maximum acuity may generate answers that seem sensible, but that often need to be examined and questioned carefully.

Stakeholders: Algorithms affect stakeholders directly or indirectly. Often different stakeholders have different or even conflicting interests.

Embedding. Algorithms may be embedded within other algorithms. For example, a decades-old optimization algorithm might be embedded within a situation-specific algorithm for assigning orders to available trucks in a specific setting.

Fitness for purpose. An algorithm's fitness for purpose is determined by 1) the form, operation, and goals of work systems within which it operates, 2) its impact on human and nonhuman actors within the work system, and 3) its impact on other stakeholders such as recipients or users of whatever the work system produces.

4 The Work System Perspective and Work System Theory

The work system perspective (WSP) is a general approach to understanding systems in organizations based on treating those systems as work systems. The core of WSP is work system theory (WST), which consists of three components: the definition of WS plus two frameworks (Figure 1) for understanding a work system: the work system framework (a static view for summarizing how a work system operates) and the work system life cycle model (WSLC - how a work system evolves through planned and unplanned change). Earlier confusion about the relationship between core ideas and various extensions of the work system approach was clarified when [2] identified WST as a core of ideas underlying the work system method. WSM had been developed over several decades as a semi-formal systems analysis method for business professionals. Various versions of WSM have been tailored to instructional needs of different courses, most of which were directed employed MBA and Executive MBA students. Individual students or teams of students used WSM templates to produce over 700 management briefings recommending improvements of problematic IT-reliant WS during 2003-2017, mostly in their own organizations (e.g., [3,4]).

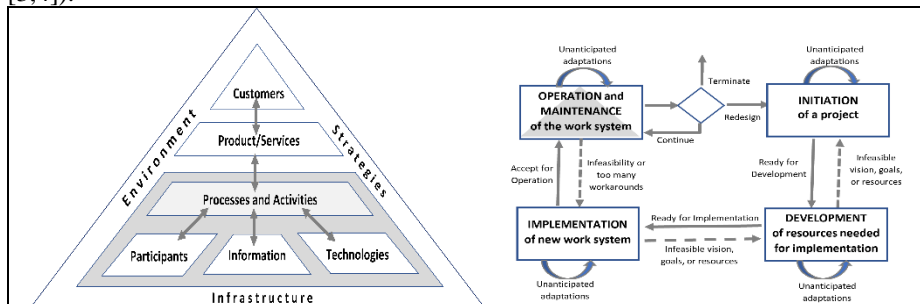


Fig. 1. Work system framework and work system life cycle model

Definition of WS. A work system is a system in which human participants *and/or* machines perform work (processes and activities) using information, technology, and other resources to produce specific product/services for internal and/or external customers. [2]. The first *and/or* addresses trends toward service-orientation and automation of work by saying that work systems may be sociotechnical (having human participants who do some of the work) or totally automated (having all of the work done by machines). A WS usually is identified based on what it is designed to accomplish and not based on software that it uses.

ISs and projects as special cases of WS. An IS is a WS most of whose activities are devoted to capturing, transmitting, storing, retrieving, deleting, manipulating,

and/or displaying information [5, 6]. An IS may be sociotechnical (e.g., financial analysts creating economic projections with the help of modeling software) or totally automated (e.g., computers generating economic projections automatically). Projects are another important special case. A project is a WS designed to produce specific product/services and then go out of existence. Software development is a type of project that can be executed in many ways.

WS framework: elements of a basic understanding of a WS. The nine elements of the WS framework outline a basic understanding of a WS's form, function, and environment during a period when it retains its identity even as incremental changes may occur, such as minor process changes, personnel substitutions, or technology upgrades. *Processes and activities, participants, information, and technologies* are completely within the WS. *Customers and product/services* may be partially inside and partially outside because customers often participate in activities within a WS and because product/services take shape within a WS. *Environment, infrastructure, and strategies* are outside of the WS even though they have direct effects within a WS.

Work system life cycle model (WSLC): how work systems change over time. WSs (including ISs) evolve through a combination of planned change and unplanned change via adaptations and workarounds. Significant changes typically affect multiple elements of the work system framework, not just technology. Planned change projects include initiation, development, and implementation phases. That separation between the phases is more in the spirit of waterfall logic and less in the spirit of today's agile approaches and trends in the direction of DevOps. Many topics in the WSLC remain valid, however, such as the emphasis on work system changes rather than just software development, the emphasis on evolution over time rather than one-time projects, the simultaneous importance of planned and unplanned change, and the relevance of sub-phases within each phase even when several phases are merged.

5 Example Illustrating the Work System Perspective on AI

Table 2 is a work system snapshot (a tool from WSM) of a hypothetical hiring system that is used to illustrate a work system perspective in a situation that might involve AI. In this example, PQR Corp implemented a new hiring work system two years ago to improve a previous hiring work system that absorbed too much effort inside PQR Corp and was so slow that good candidates sometimes went other companies before receiving offers. Also, it hired too many candidates who proved unsuitable and left before becoming productive. The new hiring work system used AlgoComm and AlgoRank from a suite of software tools provided by AlgoCorp. AlgoComm provided capabilities for posting job ads, receiving applications, setting up interview appointments, and performing other communication with candidates. AlgoRank ranked candidates based on job criteria and a neural network application driven by AlgoCorp's extensive database of job qualifications, salaries, and other information. Both AlgoComm and AlgoRank are algorithms that perform specified processing. AlgoRank can be seen as an AI application, whereas AlgoComm seems more like typical information processing even though certain parts of it apply AI technologies such as natu-

ral language processing (NLP). After two years of experience, management is once again dissatisfied with the hiring work system. Excessive effort and delays have been reduced, but interviewers and applicants find the AlgoComm interface mechanical, uninviting, and lacking a human feel. Also, the last six months saw three disastrous hires despite use of AlgoRank capabilities. Management wants to launch a new project to upgrade the hiring work system once again.

Table 2. Work System Snapshot of the Current Hiring System

Customers		Product/services	
<ul style="list-style-type: none"> • Applicants • Hiring manager • Larger organization • HR manager (who will use the applications to analyze the nature of applicants) 		<ul style="list-style-type: none"> • Applications (which may be used for subsequent analysis) • Job offers • Rejection letters • Hiring of the applicant 	
Major activities and processes			
<ul style="list-style-type: none"> • AlgoComm publicizes the position. • Applicants submit resumes to AlgoComm. • AlgoRank selects shortlisted applicants and sends the list to the hiring manager. • Hiring manager decides who to interview. • AlgoComm sets up interviews. 		<ul style="list-style-type: none"> • Interviewers perform interviews and provide comments about applicants. • AlgoRank evaluates candidates. • Hiring manager makes hiring decision. • AlgoComm notifies applicants. • Applicant accepts or rejects job offer. 	
Participants	Information		Technology
<ul style="list-style-type: none"> • Hiring manager • Applicants • Other employees who perform interviews 	<ul style="list-style-type: none"> • Job requisition • Job description • Advertisements • Job applications • Cover letters • Applicant resumes 	<ul style="list-style-type: none"> • Applicant short list • Information and impressions from the interviews • Job offers • Rejection letters 	<ul style="list-style-type: none"> • AlgoComm • AlgoRank • Office software • Internet

Interpretation Based on WST. The work system snapshot in Table 2 summarizes the work system, which involves much more than AlgoComm and AlgoRank. The hiring work system uses AI, but should not be viewed as an AI system. The transition from the previous system to the current system started with an *initiation phase* in the WSLC (Fig. 1) in which management decided to develop a new hiring system using software from a vendor. The *development phase* of the WSLC involved acquiring resources needed for implementation in the organization. In this instance, developers selected AlgoCorp as a vendor, installed AlgoCorp's software, set a group of parameters to fit it to PQR Corp's needs, and adapted AlgoCorp's training material for PQR Corp's users. In the *implementation phase* (in the organization was quick because of the relatively simple nature of the hiring process. The subsequent *operation and maintenance phase* continued for two years during which AlgoCorp updated the neural network component of AlgoRank automatically once a month to reflect changes in the job market. During those two years several incidents occurred where managers

worked around the standard process (enacting what the WSLC calls *unanticipated adaptations*) when especially talented individuals became available and might have been hired by a competitor. A division VP was consulted in one case but learned about the other workaround months later. Management is looking now at a new *initiation phase* to launch a project that will improve the hiring system further.

Using Facets of Work to Look More Deeply. The situation can be observed more deeply by using the idea of *facets of work*, which is an extension of WST. As explained in two recent conference papers (blinded), the idea of facets of work grew out of research attempting to bring richer and more evocative concepts to systems analysis and design (SA&D) to expand its scope and to facilitate analyst/stakeholder interactions. The notion of facet is an analogy to how a cut diamond is a single thing that has multiple facets. Table 3 identifies 18 facets of work, all of which describe a unique aspect of the activities that occur. Table 3 briefly mentions issues that many facets of work highlight for thinking more deeply about how algorithms associated with AI might help in generating better results. The 18 facets in Table 3 were chosen because they are easily understood, widely applicable, and associated with concepts and knowledge related to business situations. All 18 facets satisfy a series of criteria related to usefulness related to a very wide range of systems in organizations: They apply to both sociotechnical systems and totally automated systems; they are associated with many concepts that can be used for analyzing system-related situations; they are associated with evaluation criteria and typical trade-offs; they have sub-facets that can be discussed; they bring open-ended questions that are useful for starting conversations. Other researchers might have used some other number of facets that satisfy those criteria. Also, facets do not have to be independent, e.g. how decision-making often involves communication. The main point for current purposes is that each facet provides a lens for thinking about a work system that uses algorithms (or AI).

Table 3. Issues related to potential use of algorithms (specifically AI) in the hiring system

Facet	Issues related to potential use of AI in the hiring system
Making decisions	How could AI support decisions more fully in this system? Should AI suggest decisions or make decisions?
Communicating	How can AI explain how it makes or suggests decisions? How can AI help work system participants communicate more effectively?
Processing information	Can AI play any special role in capturing, transmitting, storing, retrieving, deleting, manipulating, or displaying information?
Thinking	Are there any areas in which it would be beneficial for AI to replace or augment thinking done by work system participants?
Representing reality	Does AI represent reality in a biased way? For example, what about possible bias or omissions in the dataset used to train the neural network?
Providing information	Could AI provide more meaningful information to work system participants than would otherwise be available?
Applying knowledge	Could AI identify and provide specific knowledge that would help in evaluating applicants?
Planning	Could AI help in planning hiring and interview schedules in ways that minimize interference with ongoing work by interviewers?
Controlling	Could AI help in controlling interview processes to assure that the strong-

execution	est candidates received priority in screening processes?
Improvising	How could AI support any necessary improvising, such as identifying appropriate workarounds when standard processes prove cumbersome?
Coordinating	How could AI support better coordination between interviewers and between applicants and interviewers in terms of convenience?
Performing physical work	(This is not relevant for a hiring system.)
Performing support work	How could AI perform support work that might help interviews proceed more efficiently and with fewer interruptions?
Interacting socially	Could AI do more to support social interactions during interviews and evaluation processes?
Providing service	How could AI make the entire hiring experience seem like more of a service to applicants (thereby swaying them in the firm's favor)?
Creating value	How could AI help both applicants and interviewers feel that they receive more value for their efforts across the hiring process?
Co-creating value	How could AI support active value co-creation by applicants and interviewers and by interviewers and hiring managers?
Maintaining security	How could AI help in maintaining information security for applicants, interviewers, and the firm as a whole?

Using dimensions of smartness to look more deeply at AI applications. As noted by (Iansiti and Lakhani, 2020) and many others, existing AI applications can be viewed as “weak AI” because they address individual, highly constrained problems such as those in Table 1 and in the hiring example. That approach to AI has generated important breakthroughs and efficiencies in many situations, but it is nothing like science fiction dreams of “strong AI” exhibited by humanoid robots that can reason and interact at a human or superhuman level [7]. The dimensions of smartness in devices and systems can be used to visualize current AI capabilities.

Recent papers and literature reviews have been devoted to the idea of smartness in the context of things, devices, systems, cities, and so on (e.g., [8, 9, 10, 11,12, 13, 14]). This paper uses a classification matrix for smart capabilities organized around four categories: information processing, internal regulation, action in the world, and knowledge acquisition [8]. Each category includes a set of separate capabilities, in essence separate dimensions that can be described on a continuum from not smart to somewhat smart to extremely smart based on a complex definition of smart: “Purposefully designed entity X is smart to the extent to which it performs and controls functions that attempt to produce useful results by applying automated capabilities and other physical, informational, technical, and intellectual resources for processing information, interpreting information, and/or learning from information that may or may not be specified by its designers.” Table 4 identifies 23 separate dimensions of smartness, each associated with one of four categories. Every dimension in all four categories in Table 4 is a continuous variable extending from not smart at all to extremely smart. Each dimension potentially covers a range that goes from not smart at all to increasing levels of smartness including scripted execution, formulaic adaptation, creative adaptation, and unscripted or partially scripted invention. [8]. Very few existing systems achieve anything close to the higher levels of smartness.

Table 4. Dimensions of smartness related to four categories of smartness [8]

Category of smartness	Dimensions of smartness
Information processing	capture, transmit, store, retrieve, delete, manipulate, display information
Internal regulation	self-detection, self-monitoring, self-diagnosis, self-correction, self-organization
Action in the world	sensing, actuation, coordination, communication, control, physical action
Knowledge acquisition	sensing or discovering, classifying, compiling, inferring or extrapolating from example, inferring or extrapolating from abstractions, testing and evaluating

The categories and dimensions in Table 4 can be used to see that the hypothetical hiring example is not very smart even though it uses AI. AlgoComm processes information by using mechanical, pre-specified capabilities when it captures, transmits, stores, retrieves, deletes, manipulates, and displays information. The neural network that provides AlgoRank’s parameters for ranking candidates performs a type of knowledge acquisition (classifying and compiling) using techniques that are best described as scripted execution. Neither AlgoComm nor AlgoRank demonstrate internal regulation or action in the world.

6 Application of WST, Facets of Work and Dimensions of Smartness to Real Examples

The hypothetical hiring system example was designed to illustrate the relevance of WST, facets of work, and dimensions of smartness in describing AI applications. This section applies those ideas to five real world AI applications mentioned at the “MIT AI Conference 2020: AI for a Better World” presented in a virtual format by the MIT Club of Northern California during three hours per day during July 14-18, 2020. Three were discussed in 25-minute presentations that covered a variety of business, personal, and AI topics; one was a 10-minute summary presented in a research slam; one was discussed in an interview that covered many other topics. All involved applications that might be described as digitalization, including at least one algorithm that most observers would associate with AI. The topic of every presentation could have been documented in a full paper if the presenters had been available for in-depth interviews. Since they were not available, descriptions in this paper are interpretations of their comments in the webinars. That suffices for current purposes, where the goal is to demonstrate the relevance of WST, facets of work, and dimensions of smartness for characterizing real world AI applications and their impacts.

Five applications of AI will be discussed. For the sake of uniformity, they are identified as A1, A2, etc., and will be named based on its purpose: (A1) detecting fetal heart defects, (A2) identifying defective manufactured parts, (A3) helping people learn course-related content, (A4) receiving and responding to typical IT help requests, and

(A5) generating organized notes from meetings. All of those applications are designed to be integrated into a current or future work system. The information processing (i.e., capturing, storing, retrieving, deleting, manipulating, and/or displaying information) in each case involves at least one algorithm that is associated with AI. All are or are meant to become commercial products that need to be adjusted to the application context in a variety of ways. They use different forms of supervised and/or unsupervised machine learning, with some using small training datasets initially while others require much more data before producing useful results.

The brief discussion of each application includes the nature of the work system (Figure 1, Table 2) that is being supported, the goal of the AI application within that work system, brief comments about how the AI application was developed (WSLC in Figure 1) and how it was or can be implemented as part of a work system, the main facets of work (Table 3) that are supported or automated, and a brief comment about where the application stands in relation to dimensions of smartness (Table 4).

(A1) Performing fetal screening for heart defects. [15] Congenital heart disease (CHD) is the most common birth defect, but it is still rare, occurring in only 1% of live births. Fetal ultrasound screening is recommended at 20 weeks of gestation and should be able to detect over 90% of CHD. Unfortunately, the actual rate of both true positives (sensitivity) and true negatives (specificity) is less than 50% because the rarity of CHD makes it difficult to build and maintain skills and because of difficulties of ultrasound imaging. Researchers created A1 as an ensemble of neural networks to detect CHD. They used clinical guidelines (medical knowledge) to identify five key screening views of the heart instead of looking at thousands of images for each ultrasound. This allowed them to perform an analysis based on a training dataset of 100,000 images from 1,300 ultrasounds. Later, they tested the result on a much larger dataset of 4.4 million images from 4,500 ultrasounds. Their model achieved 95% sensitivity and 96% specificity, far better than the levels in current practice. Their research included several steps. They trained a convolutional neural network to distinguish those five views. They deployed diagnostic classifiers to determine whether a fetal heart was normal for each view. They combined those classifiers to create a combined diagnostic score indicating whether the heart was normal. Overall, they found that deep learning combined with clinical knowledge and expert annotation of cases resulted in a possibility of an important way to improve on current practice.

In the future, A1 might be used by physicians during exams as they try to decide whether a 20-week fetus has CHD. In terms of facets of work, A1 will perform extensive *processing of information* to *represent reality* with a diagnostic score that *provides information* to physicians that will help them in *making decisions* about leading to treatments. In terms of dimensions of smartness, A1 *processes information* in a prescribed manner to create useful diagnostic scores. A1 does not exercise *internal regulation* or *take action* in the world. A1 does not *create knowledge*. The researchers

previously used neural networks to perform research to create the knowledge that is encoded in the new diagnostic system.

(A2) Finding defects in electronics manufacturing. [16] Inefficiencies in electronics manufacturing waste 20%-30% of expenses through scrap and rework, product returns, mistakes and experiments, and underutilized human resources. Methods for monitoring problems include automated optical inspection, functional testing, daily build reports, failure analysis reports, and analysis of customer returns. Underlying issues often are dark yield problems, i.e., defects that cannot be found through a test of function but that may cause a unit to fail later – incorrect cable routing, cold solder failures, glue overflow, connectors not fully mated, and misassembled parts such as a screw that is not fully inserted. With the COVID pandemic, engineers are prevented from going to remote factories. Merely taking pictures of work in progress or completed units in the factory is insufficient because solving problems requires tracking specific units back to specific production steps where their problems occurred. Trying to record a complete history of production units including photos during production generates a great deal of data that might have to be transmitted to the cloud from factories in remote locations where data transmission capabilities are limited.

A2 compares production units to other production units at a specific point in assembly. It can start with as few as 30 initial units before production stabilizes. It identifies anomalies such as tape and label defects, missing foams that keep components in place, missing functional parts, incorrect cable routing, and glue issues. A2 reduces the amount of labeling of defects that people need to do by sorting the images in order from totally conforming to highly nonconforming, thereby helping with decisions on cutoffs for labeling defects as consequential or not. Users can identify areas where problems are likely to occur, but that is not necessary. Thus, A2's algorithms can be trained without examples that are labeled in advance as defects. All of the training is done using software in the cloud, not on premises.

A2's work system is manufacturing of electronic items such as phones. Its algorithms are used for quick identification of defective units, even with dark yield problems, before additional defective units are manufactured. The training uses images of important parts (e.g., the front of a phone) and identifying anomalies that differentiate one unit from others. A person decides whether a unit's anomaly is serious enough to declare the unit defective. After training, algorithms can be used to monitor production to find defective units. The relevant facets of work touched directly by A2 are *processing information* (capturing images and identifying anomalies between units), *providing information* about anomalies, and *making decisions* by identifying defective units after A2 has been trained. In terms of smartness, A2 *processes information* using pre-defined scripts. It does not *perform self-regulation* or *take action in the world*. It acquires knowledge through the training process. The presentation implied that training on specific problem areas can be repeated if the produce design changes

(A3) Supporting personalized learning in coursework. [17, 18] This research involved working with students using coursework available through the Khan Academy, which provided an anonymized dataset of 50K elementary through high school

students solving 1.4 million math problems. The dataset included a history of every problem that each student tried to solve, whether the answer was correct, and how long the student worked on the problem. The researchers trained a neural network to take as input the complete history of a student's (correct and incorrect) answers to questions and to try to predict their answer to the next question. This created a complex vector for any student that predicts whether that student would solve any particular math problem that might appear next. In aggregate across all of the students, the neural network learned which skills are needed to answer any question. It represented the pedagogical structure of the mathematics that the students were trying to learn. That knowledge can lead to hints about what other problems students will be able to solve after they acquire a specific skill. In effect this knowledge graph describes in a data-driven way the prerequisite structure in learning mathematics and therefore provides a data-driven window into the learning process. A3 was developed in research reported in [18]. Real world applications are easy to imagine although the closest the webinar discussion [17] came to discussing actual applications was a few comments about a MOOC that taught coding. The very large set of student exercises and related comments by teaching assistants might be a step toward automation of some aspects of grading of coding exercises.

In a real-world application of something like A3, the work system would be students trying to learn specific coursework. A3 or something like it would hasten learning by predicting immediate difficulties students might experience and by looking ahead to provide an optimal learning sequence. In terms of facets of work, the learning management system would *process information* by storing and retrieving the student's history. It would use a student's history and a course-related neural network to *decide* what problems the student should see next. It would *represent reality* as the student's progress to date. A learning management system would *communicate* with the student through online interactions. A3 would *control* the learning process to maximize learning. In terms of smartness, all of the *information processing* would be done based on scripts that use the current state of the recurrent neural network. A3 would perform *internal regulation* and *action in the world* in the sense of using each student's history and the structure of the subject matter to decide what the student should see next. It would *acquire knowledge* by applying pre-specified neural network techniques to deepen its own knowledge as students answers problems.

(A4) Receiving and responding to IT help desk requests. [19] The firm Moveworks provides a chatbot for handling IT help requests for firms. The average time before an agent looks at an IT help request averages 5 hours and a response often takes 3 days. This problem cannot be solved with big data approaches that work in the consumer space (e.g. billions of items from webpages, documents, etc) because the IT help desk of a firm with 1000 employees might have 100 laptop requests and 10 VPN connection requests every year, not enough to serve as a training dataset for deep learning related to IT requests in that firm. As a result, chatbots often rely on hard coded logic that leads to frustrating endless loops for users. One of the problems with IT requests is that the requests are often ambiguous, e.g., "how do I get my laptop running?" The A4 approach was to build a conversational AI system that uses ma-

chine learning with “small data.” The trick in teaching the neural network was to abstract from sentence data by labeling recurrent elements, i.e., using labels that describe categories rather than instances (e.g., PC rather than Dell vs. HP vs. Lenovo). Converting sentences such as “I need Trello capability” or “Joe needs a Windows password” into a more general form like “\$PERSON wants \$SOFTWARE access” was a starting point for generating a large number of possible sentences that can be linked to actual help desk requests. A4 uses “collective learning” by applying the same learning approach across many firms that have IT help desks. It also used “transfer learning” by extracting universal language patterns (e.g., that good answers often have instructions in the form of lists) that can be applied across domains. Initially they used stackoverflow.com, a website for software developers that contains millions of help requests and related answers. The ultimate result is a chatbot that can completely answer around 40% of help requests and can escalate the others to human operators. As a result, the human operators handle many fewer IT tickets, a great saving in the use of a scarce resource.

A4 is part of a WS that answers IT help requests. Facets of work include *making decisions, communicating, processing information, providing information, and representing reality*. In terms of smartness, every type of *information processing* is present in a scripted form. A4 performs *internal regulation* by recognizing the current state of its dialogue with a user and trying to respond appropriately. It *takes action in the world* by engaging in a dialogue with users. It *acquires knowledge* from its usage.

(A5) Producing useful notes from meetings. [20] Fireflies.ai is a commercial product that records meetings and generates transcripts automatically. A5 is the basis of “Fred,” an automated voice assistant that records meetings, produces transcripts, and performs other tasks to make meetings and their aftermath more efficient. A user invites fred@fireflies.ai to an online meeting on Zoom, WebEx, or other platforms. Fred captures and transcribes voice conversations, indexes the notes to make them useful, and routes the notes to anyone who should receive them. Action items can be transferred automatically to project management systems such as Trello or to customer relationship management systems such as Salesforce without doing a lot of manual work. Maintaining a complete history of meetings makes it possible to find details of meetings that may have happened months ago.

In effect, A5 is the technical basis of an automated work system that is created through three main steps, two of which are in the WSLC *development phase*. The first step is collecting relevant language data (sentences, keywords, etc.) from users and public sources, storing the data, and labeling the data to make it useful. For example, the founders of Fireflies.ai developed some of their ideas by labeling 20K data points (basically sentences) from their own meeting notes and recordings by using a yes/no binary classification model (important or unimportant). The initial trial use of the resulting model led to more language data that could be incorporated. The second main step was to get the model running using available software that fit the model and making sure that errors and duplicates in the data (terms and sample sentences) were eliminated. The third step was benchmarking and performance improvement (e.g., minimizing false positives and false negatives) as part of an *operation and mainte-*

nance phase. Subsequent use in different industries is accommodated by having users pick a domain such as health care or sales when they sign up. Each industry has keywords that appear on user dashboards. Users add other keywords. Meeting agendas and user edits to transcripts provide more industry-specific jargon.

The relevant facets of work *processing information* (recording and transcribing meetings), *providing information* in the form of transcripts. In terms of smartness, A2 *processes information* using pre-defined scripts. It does not *perform self-regulation*. It *takes action in the world* by producing and distributing transcripts. It acquires knowledge by improving its language models every time it is used.

7 Conclusion

This paper’s contribution is way to describe and characterize AI applications and their impacts instead of just talking about AI in general giving cherry-picked examples. This paper used a hypothetical hiring example to illustrate the ideas and then applied those ideas to five real examples. The five examples illustrate a number of points that are not evident from many attempts to talk about AI in general.

Integration with work systems. (Figure 1, Table 2). In all five cases, the AI algorithm was part of an actual work system (A2, A4, A5) or was developed as research with a high potential for application in work systems (A1, A3).

Facets of work. (Table 3) All of the examples *process information*, *provide information*, and *represent reality*. All *create value*. Algorithms in A1, A2, and A3 contribute to *making decisions* that matter. A4 and A5 *perform support work*. A2 helps in *controlling execution* by identifying anomalies. The chatbot in A4 *communicates* with customers of the help desk. A1 and A2 *communicate* in a more structured way.

Smartness (Table 4) All five examples *process information* in a scripted way. None exhibit *internal regulation*. A1 and A2 identify problems but do not *take action in the world*. A4, and A5 take action in the world, and A3 has a potential to do so. All use knowledge that was built into neural networks, but only A3 *acquires knowledge*. The relevant dimensions of smartness in all cases are handled through scripted execution of algorithms rather than by autonomous modification of algorithms.

Importance of domain knowledge. The neural networks in A1, A2, A4, and A5 all depended on domain knowledge built into their design. A3 is more like unsupervised learning based on paths that different students took in learning subject matter. The dream of unsupervised learning seems relevant mostly for situations where situational knowledge is not necessary, such as the way Open AI’s widely touted GPT-3 system [21] generated 1.5 billion “parameters” that describe the likelihood that specific words will occur next to other words in a vast training dataset of texts.

Big data or little data. Some AI algorithms such as GPT-3 are built on huge data bases, but A2 could start being useful after training with only 30 examples. A1, A4, and A5 also used knowledge as a way to reduce the size of training datasets.

Visibility to users. General discussions of AI often mention the lack of visibility to users. A1 is based on ultrasound data that is understood by highly skilled users. A2 finds anomalies that are visible. A3 is hidden within a learning management system. A4 and A5 perform support work where most errors are easily identified.

Is AI inherently ethically suspect? General discussions of AI frequently focus on harm that can occur through unethical uses of AI. At this point it is obvious that certain types of AI applications such those that identify people or suggest important decisions related to specific individuals are susceptible to many types of intentional or unintentional bias. Such bias may occur in the training datasets (e.g., inadequate coverage of a specific gender, age, or ethnic group) and/or in the logic of the work system (e.g., maximizing profit in a way that disadvantages specific groups, such as people who live in specific areas or have specific characteristics). On the other hand, it is also obvious that almost any technology can contribute to work systems that harm people. The issue is not with AI as a category, but rather with work systems and/or algorithms that fail to consider fair treatment for all customers and other stakeholders.

References

1. Schultze, U., Aanestad, M., Mähring, M., Østerlund, C., Riemer, K.: Living with Monsters? Social Implications of Algorithmic Phenomena, Hybrid Agency, and the Performativity of Technology. Proceedings of IFIP WG 8.2 Working Conference on the Interaction of Information Systems and the Organization. December 11–12. San Francisco, CA, USA. (2018).
2. Alter, S.: Work system theory: overview of core concepts, extensions, and challenges for the future. *Journal of the Association for Information Systems*, 14(2), pp. 72–121. (2013).
3. Truex, D., Alter, S., Long, C.: Systems analysis for everyone else: Empowering business professionals through a systems analysis method that fits their needs. In: Proceedings of the European Conference on Information Systems (2010).
4. Truex, D., Lakew, N., Alter, S., Sarkar, S.: Extending a systems analysis method for business professionals. In *European Design Science Symposium* (pp. 15-26). Springer, Berlin, Heidelberg. (2011).
5. Alter, S.: The work system method: Connecting people, processes, and IT for business results. Work System Press, Larkspur, USA (2006).
6. Alter, S.: Defining information systems as work systems: implications for the IS field, *European Journal of Information Systems*, (17) pp. 448-469, (2008).
7. Iansiti, M., Lakhani, K.R.: *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*. Harvard Business Press, Cambridge, USA, (2020).
8. Alter, S.: Making Sense of Smartness in the Context of Smart Devices and Smart Systems, *Inf. Syst. Frontiers*, 22(2), pp.275-280
9. Strozzi, F., Colicchia, C., Creazza, A., Noè, C.: Literature review on the ‘Smart Factory’ concept using bibliometric tools. *International Journal of Production Research*, 55(22), pp.6572-6591 (2017).
10. Romero, M., Guédria, W., Panetto, H., Barafort, B.: Towards a characterisation of smart systems: A systematic literature review. *Computers in Industry*, 120, in press, (2020).

11. Elbanna, A.R., Dwivedi, Y., Bunker, D., Wastell, D.: The Search for Smartness in Working, Living and Organising: Beyond the 'Technomagic'. *Inf. Syst. Frontiers*, 22(2), pp.275-280. (2020)
12. Langley, D.J., van Doorn, J., Ng, I.C., Stieglitz, S., Lazovik, A., Boonstra, A.: The Internet of Everything: Smart things and their impact on business models. *Journal of Business Research*, in press. (2020).
13. Wolf, V., Bartelheimer, C., Beverungen, D.: Establishing Smart Service Systems is a Challenge: A Case Study on Pitfalls and Implications. *Wirtschaftsinformatik 2020 Zentrale Tracks*. (2020).
14. Anke, J., Poepplbuss, J. and Alt, R., 2020.: Joining Forces: Understanding Organizational Roles in Inter-organizational Smart Service Systems Engineering. *Wirtschaftsinformatik 2020 Zentrale Tracks*. (2020).
15. Arnaout, R.: Improving fetal heart screening with deep learning, <https://www.youtube.com/watch?v=76WvsHhCyko>, (2020), last accessed 2020/8/1
16. Shedietsky, A-K.: How COVID-19 is driving adoption of AI in electronics manufacturing. <https://www.youtube.com/watch?v=DiuuVTJKoLs>, (2020), last accessed 2020/8/1
17. Ganguli, S.: Learning in Humans and Machines: From Optimizing Our Biology to Colluding With AI's, <https://www.youtube.com/watch?v=OL89f0mnJ00>, (2020), last accessed 2020/8/1
18. Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L. J., Sohl-Dickstein, J.: Deep knowledge tracing. In *Advances in neural information processing systems* (pp. 505-513). (2015).
19. Shah, B.: How Conversational AI Will Empower One Billion Knowledge Workers, https://www.youtube.com/watch?v=bzBZA_2kkNE, (2020), last accessed 2020/8/1
20. Udotong, S.: Practical Deep Learning: How It Is Used in Production and Already Revolutionizing Enterprise Work. <https://www.youtube.com/watch?v=ywhY2DCIZqU> (2020), last accessed 2020/8/1.
21. Manjoo, F.: How Do You Know a Human Wrote This? *New York Times*, Editorial, July 29, 2020. <https://www.nytimes.com/2020/07/29/opinion/gpt-3-ai-automation.html>, last accessed 2020/8/1.