DATA-DRIVEN DIALOGUE pdf

1: Data-Driven Conversation - Microsoft Research

Data Driven Dialogue Developed by the Teacher Development Group, Based on work presented by Nancy Love, author of "Using Data/Getting Results,"

Yes, this page paper does include a catalogue of data sets with dialogues from different domains, but it also includes a high level survey of techniques that are used in building dialogue systems aka chatbots. In particular, it focuses on data-driven systems, i. This particular paper is focused on corpus-based learning where you have been able to build up, or have access to, a data set on which you can train your models. Out of scope is training dialogue systems through live interaction with humans â€" but there are some references to follow on this so I may well return to that topic later on in this mini-series. Anatomy of a dialogue system The standard architecture for a dialogue system looks like this: Natural language interpretation and generation are core NLP problems with applications well beyond dialogue systems. For building chatbots, where we assume written input and output, the speech recogniser and synthesiser can be left out. I had naively assumed that if you had a good working system that can deal with textual inputs and outputs, it would be a simple matter of bolting a speech-to-text recogniser in front of the system in order to build a voice-driven assistant. The distinction between spoken and written dialogues is important, since the distribution of utterances changes dramatically according to the nature of the interactionâ€¦. Spoken dialogues tend to be more colloquial, use shorter words and phrases and are generally less well-formed, as the user is speaking in a train-of-thought manner. Conversely, in written communication, users have the ability to reflect on what they are writing before they send a message. Written dialogues can also contain spelling errors or abbreviations, which are generally not transcribed in spoken dialogues. Even written dialogue â€" for example for movies and plays, and in fictional novels â€" has apparent distinctions from real speech. Which leads to this wonderful observation: Within dialogue systems we can distinguish between goal driven systems â€" such as travel assistants or technical support services â€" where the aim is to accomplish some goal or task, and non-goal driven systems such as language learning tools or computer game characters. Most startups building chatbots will be building goal driven systems. Initial work on goal driven dialogue systems primarily used rule-based systemsâ€¦ with the distinction that machine learning techniques have been heavily used to classify the intention or need of the user, as well as to bridge the gap between text and speech. Research in this area started to take off during the mid 90s, when researchers began to formulate dialogue as a sequential decision making problem based on Markov decision processes. Commercial systems to date are highly domain specific and heavily based on hand-crafted features. For example, to predict the intent of a user in a dialogue, conditioned on what they have said. Here the intent is the label, and the conditioned utterances are called conditioning variables or inputs. Discriminative models can be similarly applied in all parts of the dialogue system, including speech recognition, natural language understanding, state tracking, and response selection. One popular approach is to learn a probabilistic model of the labels, another is to use maximum margin classifiers such as support vector machines. Answering back When it comes to choosing what your chatbot is going to say e. The simpler approach is to select deterministically from a fixed set of possible responses which may of course use parameter substitution: The model maps the output of the dialogue tracker or natural language understanding modules together with the dialogue history e. This approach effectively bypasses the natural language generation part of the system. The fixed responses may have been crafted up-front by the system designers, but there are also systems that effectively search through a database of dialogues and pick the responses from there that have the most similar context: The optimal response is then found by projecting all potential responses into the same Euclidean space, and the response closest to the desirable response region is selected. More complex chatbots generate their own responses. Using a method known as beam-search they can generate highly probably responses. The approach is similar to that used in the sequence-to-sequence machine translation paper we looked at recently. Short single request-response conversations are simpler than those that need to be able to handle multiple interactive turns. For example, an interactive system might require steps of clarification from the user before being able to offer pertinent information. Indeed, this is a common scenario: This phase is sometimes referred to as the grounding process. To tackle such behaviors, it is crucial to have access to dialogue corpora with long interactions, which include clarifications and confirmations which are ubiquitous in human conversations. The need for such long-term interactions is confirmed by recent empirical results, which show that longer interactions help generate appropriate responses. Incorporating external knowledge Chatbots may rely on more than just dialogue corpora for training. When building a goal-drive dialogue system for movies, Dodge et al. They use four different subsets of data to train models for these tasks: This could include structured information â€" such as bus or train timetable for answering questions about public transport â€" typically contained in relational databases or similar. Some good sources include: Tools include part-of-speech taggers, word category classifiers, word embedding models, named entity recognition models, semantic role labelling models, semantic similarity models, and sentiment analysis models. They created an utterance-level representation by combining the word embeddings of each word, for example, by summing the word embeddings or taking the maximum w. These utterance-level representations, together with word counts, were then given as inputs to a linear classifier to classify the dialogue acts. Thus, Forgues et al. What were you saying? Tracking the state of a conversation is a whole sub-genre of its own, which goes by the name of dialogue state tracking or DSTC. It is framed as a classification problem: For example, the system may believe with high confidence that the user has requested timetable information for the current day. DSTC model include both statistical approaches and hand-crafted systems. Although none of these models are explicitly designed to address dialogue problems, the extension by Kumar et al. In this case, the episodic memory is the same as the memory used in the traditional Memory Networks paper which is extracted from the input, while the semantic memory refers to knowledge sources that are fixed for all inputs. The model is shown to work for a variety of NLP tasks, and it is not difficult to envision an application to dialogue utterance generation where the semantic memory is the desired external knowledge source. There is strong evidence that over the next few years, dialogue research will quickly move towards large-scale data-driven model approaches, in particular in the form of end-to-end trainable systems as is the case for other language-related applications such as speech recognition, machine translation and information retrievalâ€¦ While in many domains data scarcity poses important challenges, several potential extensions, such as transfer learning and incorporation of external knowledge, may provide scalable solutions.

DATA-DRIVEN DIALOGUE pdf

2: Data-Driven Dialogue - USDData Analysis Protocol Companion Website

Data-Driven Dialogue has 8 ratings and 1 review. This timely book offers school leaders a practical toolkit for structuring and facilitating collaborativ.

Yes, this page paper does include a catalogue of data sets with dialogues from different domains, but it also includes a high level survey of techniques that are used in building dialogue systems aka chatbots. In particular, it focuses on data-driven systems, i. This particular paper is focused on corpus-based learning where you have been able to build up, or have access to, a data set on which you can train your models. Out of scope is training dialogue systems through live interaction with humans â€" but there are some references to follow on this so I may well return to that topic later on in this mini-series. Anatomy of a dialogue system The standard architecture for a dialogue system looks like this: Natural language interpretation and generation are core NLP problems with applications well beyond dialogue systems. For building chatbots, where we assume written input and output, the speech recogniser and synthesiser can be left out. I had naively assumed that if you had a good working system that can deal with textual inputs and outputs, it would be a simple matter of bolting a speech-to-text recogniser in front of the system in order to build a voice-driven assistant. The distinction between spoken and written dialogues is important, since the distribution of utterances changes dramatically according to the nature of the interactionâ€¦. Spoken dialogues tend to be more colloquial, use shorter words and phrases and are generally less well-formed, as the user is speaking in a train-of-thought manner. Conversely, in written communication, users have the ability to reflect on what they are writing before they send a message. Written dialogues can also contain spelling errors or abbreviations, which are generally not transcribed in spoken dialogues. Even written dialogue â€" for example for movies and plays, and in fictional novels â€" has apparent distinctions from real speech. Which leads to this wonderful observation: Within dialogue systems we can distinguish between goal driven systems â€" such as travel assistants or technical support services â€" where the aim is to accomplish some goal or task, and non-goal driven systems such as language learning tools or computer game characters. Most startups building chatbots will be building goal driven systems. Initial work on goal driven dialogue systems primarily used rule-based systemsâ€¦ with the distinction that machine learning techniques have been heavily used to classify the intention or need of the user, as well as to bridge the gap between text and speech. Research in this area started to take off during the mid 90s, when researchers began to formulate dialogue as a sequential decision making problem based on Markov decision processes. For example, to predict the intent of a user in a dialogue, conditioned on what they have said. Discriminative models can be similarly applied in all parts of the dialogue system, including speech recognition, natural language understanding, state tracking, and response selection. One popular approach is to learn a probabilistic model of the labels, another is to use maximum margin classifiers such as support vector machines. Answering back When it comes to choosing what your chatbot is going to say e. The simpler approach is to select deterministically from a fixed set of possible responses which may of course use parameter substitution: The model maps the output of the dialogue tracker or natural language understanding modules together with the dialogue history e. This approach effectively bypasses the natural language generation part of the system. The fixed responses may have been crafted up-front by the system designers, but there are also systems that effectively search through a database of dialogues and pick the responses from there that have the most similar context: The optimal response is then found by projecting all potential responses into the same Euclidean space, and the response closest to the desirable response region is selected. More complex chatbots generate their own responses. Short single request-response conversations are simpler than those that need to be able to handle multiple interactive turns. For example, an interactive system might require steps of clarification from the user before being able to offer pertinent information. Indeed, this is a common scenario: This phase is sometimes referred to as the grounding process. To tackle such behaviors, it is crucial to have access to dialogue corpora with long interactions, which include clarifications and confirmations which are ubiquitous in human conversations. The need for such long-term interactions is confirmed by recent empirical results, which show that longer interactions help generate appropriate responses.

DATA-DRIVEN DIALOGUE pdf

3: Data Driven Dialogue â€“ School Reform Initiative

Data-Driven Dialogue (Wellman & Lipton,) is a structured process that enables a Data Team to explore predictions, go visual, make observations, and generate inferences and questions of the data before offering solutions.

Data-Driven Dialogue A process of sense making and discovery, not decision making. Applies what is known about how learners data team members make sense of new knowledge through activating prior knowledge, using vibrant visual displays, opening up extended opportunities for exploration and discovery, and engaging in dialogue about assumptions. Four Phases Predict The best predictor of new knowledge is access to prior knowledge Participants in the dialogue activate their prior knowledge and surface assumptions that they bring to the data. Predicting accomplishes several purposes: By articulating these assumptions, they can more easily set them aside and examine data more objectively. Increase in student mobility Assumption: The school has had an increase in the percentage of low SES students Go Visual Simple, large, and colorful data displays that four to six people can look at together at once are a hallmark of the Using Data Process; large charts, fish bone diagrams, chart paper, etc. Strong tendency for participants to want to leap to an explanation and interpretation before they have fully explored what can be learned from the data. USD has a larger percentage of Hispanic students than we do teachers In , USD has a large number of students not achieving as expected From to the number of students meeting standards in mathematics, according to the KMA, has increased from The average number of students failing one or more class at Emporia Middle School has increased over time. From to the mean number students failing at least one class at Emporia Middle School has decreased by 12 students. Questions to pose to help teachers move from rough observations to refined observations: Does each statement communicate a single idea about student performance? Do the statements incorporate numbers? Do the statements focus on lust those direct and observable facts that are contained in the data, without interpretation or influence? Do the statements use relevant data concepts, such as mean, media, mode, range, or distribution? Participants tap back into their prior knowledge and assumptions to generate multiple possible explanations or implications of what they are seeing in the demographic data. The challenge is to keep the teachers and the Data Team open to multiple possible explanations and not to latch on to their first conclusions. Teachers and Data Teams must be willing to dig into the data, make observations, and ask questions. Data-Drive Dialogue is the process of sense making and discovery, not decision making. How might parents and students from diverse cultural groups in our school interpret these data? Can I stop and reframe my interpretation? What would a culturally proficient interpretation of these data be? We have a lot more kids at risk and less ability to build background outside of school. I wonder if our instructional programs are geared to meet the needs of these students? Is the standards based instructional program being taught with enough rigor? Are the students being exposed to the appropriate level of reading for the content at 3rd grade? For K-8, the number of students not making standards on the Kansas Reading Assessment appears to reach a maximum at the 5th grade. What is happening instructionally in the 5th grade standards?

4: [] A Survey of Available Corpora for Building Data-Driven Dialogue Systems

The 3 phases of data-driven dialogue assist groups in making shared meaning of data. We encourage you to use this tool with your entire school staff and/or with your school leadership team at a special meeting on data.

Communal life on the African savannahs shaped our early ancestors mentally and emotionally. Like our predecessors, we are attuned to the cues of others and are predisposed to take collective action when it serves the greater good. Yet, there is really no such thing as a group. Groups emerge from collections of individuals who make choices about how and when to participate. These boundaries are the membranes through which information and resources flow in and out of the group. At one level, group members are the bodies enclosed within the boundaries of a membrane. In the study of physical science, we learn that membranes can be permeable, semi-permeable, or impermeable to various size molecules. In the same ways, the boundaries of groups and of group members vary in permeability. The key difference is that these boundaries are not fixed physical properties. Skilled group leadership and purposeful group development open these membranes within and between people, information and insight. Creating cultures of inquiry promotes boundary shifting, knowing when to focus in and when to focus out. It is the work of seeking information, processing information, decision-making, planning and implementation. These functions are also the work of individuals operating within the membrane of the group. Leading cultures of inquiry means helping members see their parts within the whole and helping them take responsibility for regulating their personal and collective permeability to perspectives, ideas, options and actions. Group Development Leading groups that work in which there is maximum participation, productivity and satisfaction, requires attention to three arenas of group development: Productive groups learn from experience by setting goals for themselves, monitoring their performance and reflecting on their practice. Experience by itself is not a reliable teacher. By focusing only on the tasks at hand, groups may get that work done but do not expand their capacities for addressing increasingly harder or more complex tasks. Many groups operate with similar blinders missing the importance of the organizing their tasks to increase their efficiency and productivity, developing their process toolkit for supporting thinking and clear communication, or purposely building relationships within the group to develop their capacities for collaboration and strengthening professional community. In practice, one of the most imposing roadblocks on the journey to engaging collaborative inquiry has been the ways in which data are used. Accountability requirements and legislated mandates have pushed the administration of formal assessments from which data are collected and reported to external agencies, including the district office, the province, the federal government or even international bodies. There tends to be lower uses of data for self-assessment, problem solving, reflection and personal discovery. That is, data is often used to prove, rather than improve. Data is perceived as something that needs to be gathered, organized and transferred to comply with the requirements of others rather than as a tool for gaining insight into and improving professional practices. Nancy Love and her colleagues describe this pattern as the difference between teachers being submissive data-givers and becoming confident data-users. Effective data use requires school leaders who know how to shape cultural conditions that support thoughtful, informed and collaborative data-driven dialogue in their schools, and classroom teachers who have the skills and processes for participating fully. School improvement is a discipline Elmore, It requires meaningful targets, timely feedback mechanisms and collaborative interactions focused by data that reach and reflect the technical core of teaching and learning in the school. To be successful, these challenging endeavors require three critical conditions for any working group: Collaborative inquiry is about shared learning and is therefore both psychologically and socially risky. For productive collaborative inquiry, it must be safe to not know. Emotional safety is a foundation for cognitive resourcefulness. Thinking is hard work. Collaborative inquiry works best when groups share a process toolkit, clear structures and deliberate protocols for tapping the thoughts of everyone in the room not just the vocal few. Questions that cue specific cognitive processes increase the effectiveness of collaborative forums. For example, questions that direct group members to predict, compare, contrast, infer cause and effect, generalize facilitate the task of data analysis, interpretation, planning and problem solving. In steadily improving schools, teachers actively question and explore individual and collective teaching practices calibrated by both student data and shared learning and teaching standards Chenoweth, These collaborations require relationships that are resilient enough to withstand close scrutiny of professional practice. Group members need skill in cognitive conflict, that is, conflict with ideas â€" not with one another Amason, et. Attention to developing shared processes, structures and effective collaborative experiences produces relationships that can withstand difficult conversations, challenging questions and respectful, public examination of teaching practices. They are the result of planning, problem-solving and reflection-on-action on the part of both group members and group leaders. Successful group leaders see the group as it might be, not as it is. This requires a developmental lens for group development and a willingness to invest in thoughtful capacity building and not just immediate task accomplishment. A similar developmental lens is important for group leaders to embrace for themselves as well. We all need to learn to project ourselves as group leaders into the future and operate in the moment with that bigger picture in mind. A willingness to grow and develop as a group leader conveys an important message about the purposes and values of professional collaboration and the purposes and values of professional learning. An important dimension in successful management teams. Organizational Dynamics, 24 2 , How member change and continuity affect small group structure, process, and performance. Small Group Research, 24, Academic success in unexpected schools. The limits of change. Harvard Education Letter, 18, No. Creating and leading cultures of inquiry. A practical guide for school improvement in mathematics and science.

5: 4-Phase Data Dialog | Using Data for Meaningful Change

Data-Driven Dialogue: A Facilitator's Guide to Collaborative Inquiry, 2nd Edition By Bruce Wellman and Laura Lipton. This new edition offers school leaders a practical toolkit for developing powerful teams that use data for continuous school improvement.

6: A Survey of Available Corpora for Building Data-driven Dialogue Systems

3 Data-Driven Dialogue Protocol Ann R. Pearce, Ph.D. 3) Other questions to explore: "What seems to be surprising or unexpected?" 4) Decide if this is a good place to stop to gather.

7: Data Driven Dialogue

For protocol and facilitation, see Data Driven Dialogue Protocol Facilitation Plan. Protocols are most powerful and effective when used within an ongoing professional learning community and facilitated by a skilled facilitator.

8: A survey of available corpora for building data-driven dialogue systems | the morning paper

A survey of available corpora for building data-driven dialogue systems Serban et al. Bear with me, it's more interesting than it sounds:). Yes, this (page) paper does include a catalogue of data sets with dialogues from different domains, but it also includes a high level survey of techniques that are used in building dialogue systems (aka chatbots).

9: Data-Driven Dialogue, 2nd Edition: Bruce Wellman and Laura Lipton: www.enganchecubano.com: Book

Dialogues: Strategies for Conversations Pennsylvania High School Coaching Initiative May, 2 Outcomes for Our Time Together process for data driven dialogues.

DATA-DRIVEN DIALOGUE pdf

Ms word 2007 bangla tutorial Next, the coming era in science Synchronous and resonant DC/DC conversion technology, energy factor, and mathematical modeling Harrison book of internal medicine 18th edition Historical and archived parcels Skyline wilderness park Gunner Asch goes to war TOM BRGHTWIND, OR, HOW THE FAIRY BRIDGE Dry clean price list Cam Jansen and the Ghostly Mystery (Cam Jansen) Drug Enforcement Administrations alleged connection to the Pan Am flight 103 disaster The Lawn Garden Owners Manual Pirates, sugar, debtors, and Federalists : the paradoxes of antislavery political economy Origin and destiny of humanness Requesting and requiring permissions Minorities in the workplace Shareholder proposals: a platform for communication and change Asta the screwball dog : Hollywoods canine sidekick Sara Ross, James Castonguay The Paradine Case/Foreign Correspondent/The Lodger (Alfred Hitchcock Collection) Symbols for communication Adlers Multiple Percussion Solos Advanced The life of yogananda book Man as individual Geosynthetics in foundation reinforcement and erosion control systems Storymaking in education and therapy Indian attitudes towards anti-Semitism Shutting Down the System Offers in compromise : step right up! Pay pennies on the dollar! Laszlo Moholy-Nagy (Phaidon 55s) In the arms of the sky How to Write Better Resumes and Cover Letters (How to Write Better Resumes) Against the conventional wisdom Spelling for grade 1 Robert Greenes Planetomachia (1585 (Literary and Scientific Cultures of Early Modernity) Oliver and Albert, friends forever Crimes against the state Materials in world perspective Unit 5. Americas journal George Leslie Mackay, Formosas preacher and teacher. A Long Obedience Journal