Lydia Manikonda[↑], Shirin Sohrabi[‡],
Kartik Talamadupula[‡],
Biplav Srivastava[‡], Subbarao Kambhampati[↑]

[↑]Arizona State University
[‡] IBM Research

Extracting Incomplete Planning Action Models from Unstructured Social Media Data to Support Decision Making

Motivation

01

Planning and decision-making techniques depend heavily on the availability of complete and correct models.

02

Most existing planning and decision-making methods assume hand-coded, fully-specified action models *a priori*.

03

Dependence on hand-coded, fully-specified models restricts the number and types of domains where planning can be used.

“Planning” is **more than** just plan synthesis from hand-coded models

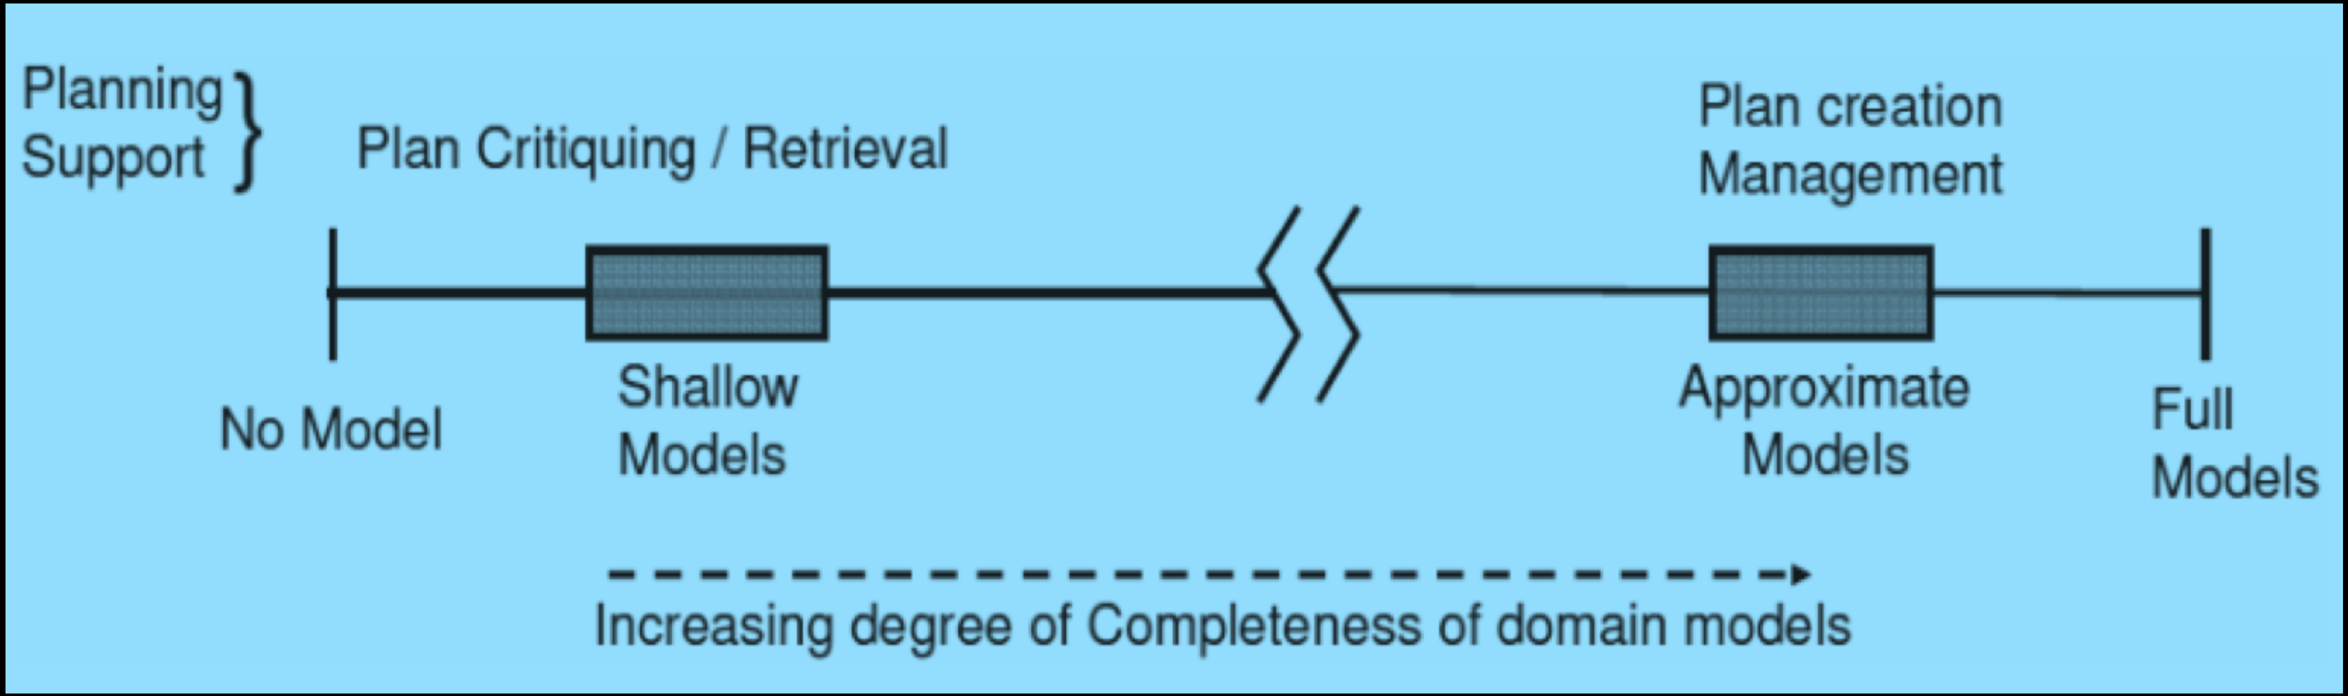

Figure 1: A Spectrum of Incomplete Domain Models and associated planning capabilities

Motivation

Planning in the Data-Centric Age
Panel on Wednesday

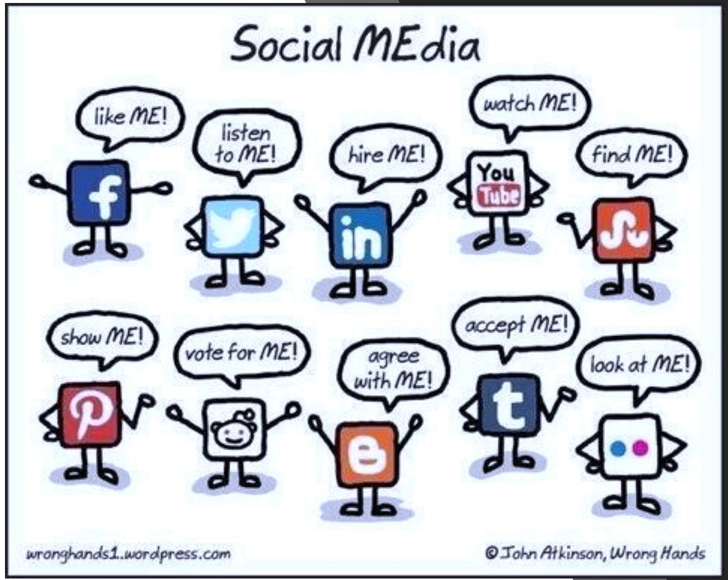

- A wealth of user-generated data is available online – especially on online social media platforms
- A lot of this data is crowd-generated wisdom about actions and plans to take in order to achieve certain high-level objectives
  - E.g. Health & Fitness, Event Planning, Travel Planning, Academic Preparation
- Growing interest in exploiting this data to provide data-based decision support, especially by extracting causal relationships
- Decision Support: Can be plan synthesis, but can also be…
  - Plan Critiquing
  - Plan Ranking
  - Plan Retrieval
  - Plan Visualization, etc.

Problem Statement

1. Can we extract approximately structured plan traces from unstructured social media posts?
2. Using these approximate plan traces, can we build an approximate causal model?

Causal models can then be utilized to automatically explain (*plan explanation or plan critiquing*), visualize, retrieve, or rank the experiences shared by users on online social media platforms.

Challenges

- *Extracting* actions from unstructured users' posts
- *Processing* the extracted actions to reduce redundancy
- *Building* plan traces from the extracted actions
- *Constructing* an action precedence graph from these traces
- *Planning* using these precedence graphs

Social Media Data: Reddit

	Quit Smoking	Couch to 5K	Wedding Planning
#Users	787	604	969
Total # of traces	1598	1131	3442
Average Trace Length	17.97	16.7	21.29
#Unique actions (orig)	1712	1299	2666
#Unique actions (model)	234	194	355
#Pre-actions	1499	1060	2795
#Post-actions	1398	982	2619

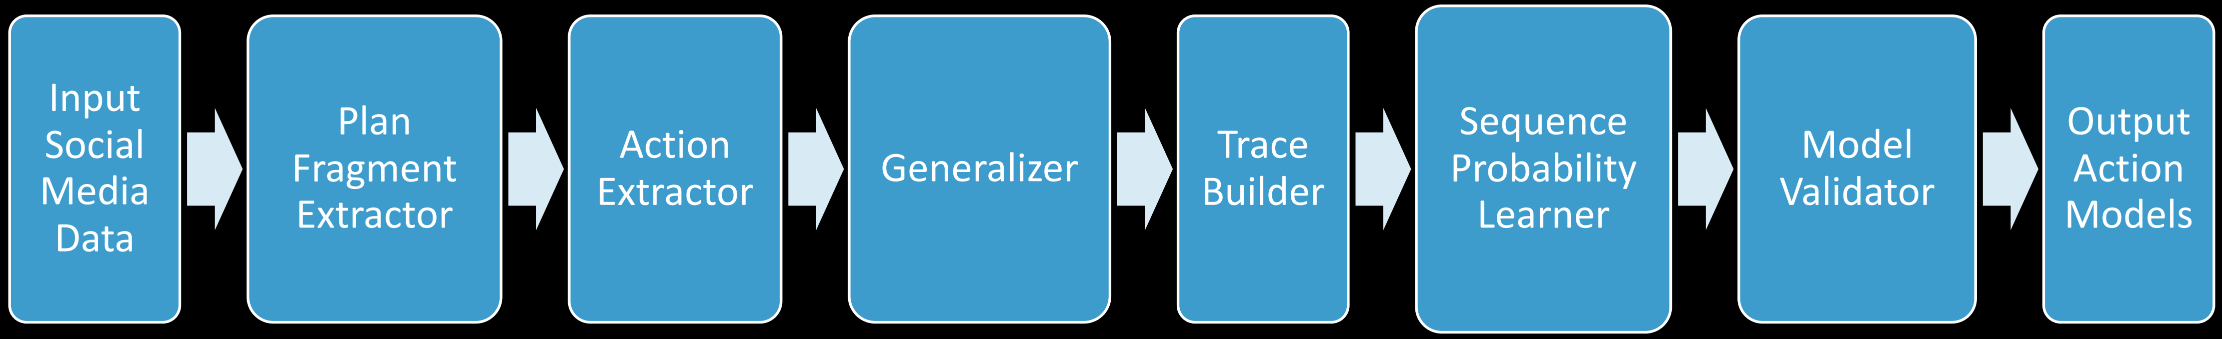

Solution Architecture

Architecture – Plan Fragment Extractor

Goal: Extract fragments from the corresponding sub-reddit

- A fragment is defined as the relevant post that contains information about achieving a given goal

→

Approach: We utilize sub-reddit tags attached to posts made by individuals on their timelines

Architecture – Action Extractor

- **Goal:** Extract action names along with their predicates
- **Approach:** Utilize the part of speech tagger to identify the verbs (actions) and nouns (predicates)

Architecture – Generalizer

- **Goal:** Due to handling unrestricted natural language text, we need to generalize a set of actions to their parent action, to remove redundancy
- **Approach:** Hierarchical clustering approach with Leacock Chodorow (LCH) similarity metric to compute the distance between any two given actions
  - LCH: Used in cases where there are distances between objects, and a taxonomy that relates them

Architecture – Trace Builder

- **Goal:** Represent the actions in the plan traces with their corresponding cluster representatives
- **Approach:** Map the corresponding cluster representatives to the actions and represent the plan traces in terms of these representatives

Architecture – Sequence Probability Learner

- **Goal:** Extract the pre-actions and post-actions for any given action by using probability metrics
- **Approach:** Data-driven conditional probability estimation between any two given actions (a_i and a_j)

$$p(a_i \mid a_j) = \frac{p(a_i \cap^* a_j)}{p(a_j)}$$

Architecture – Model Validator

- **Goal:** To measure the *explainability* of the causal model derived from the plan traces
- **Approach:** Build the model using 80% of the data (*T*) and test it using the remaining data (*T'*) by considering a metric called *explainability*:

$$T'' = T \cap T'$$

$$Explainability = \frac{|T''|}{|T'|}$$

Sample Relationships Extracted

Quit Smoking	Couch to 5K	Wedding Planning
`(:action change(ability) [:pre-action eat(gross) crave(succeed) dealt(reality)] [:post-action set(goal) run(mile) quit(smoke)] )` Candidate Explanation: Someone is craving success and is dealing with the reality of eating gross food. This person wants to change their abilities and so sets some goals, runs miles, and quits smoking.	`(:action sign(race) [:pre-action recommended(c25k) push(run) refer(program)] [:post-action begin(week) run(minute) cover(mile) know(battle) kept(pace)] )` Candidate Explanation: A person was recommended the couch to 5K and was encouraged to run. He refers to a program and so signs up for the race. He then begins to run a few minutes from the next week and covers a few miles. He knows that it is a battle but has kept the pace.	`(:action hate(dress) [:pre-action pick(dress) saw(dress) blame(problem) cost(much)] [:post-action kill(wed) find(dress) move(wedding)] )` Candidate Explanation: The person picks up the dress and sees the dress. It may cost a lot, and the person starts looking at this as a reason to kill the scheduled wedding. Eventually the person must find a new dress and move the wedding date.

Evaluation 1: Explainability (10-fold cross-validation)

- One would expect that as the clustering threshold decreases, the explainability would increase
  - **Reason:** Redundancy of actions decrease

α	Quit Smoking	Couch to 5K	Wedding Planning
2.50	65.66%	64.5%	73.39%
2.25	65.66%	64.59%	73.39%
2.0	68.41%	69.78%	77.7%
1.75	69.33%	70.67%	78.39%
1.50	80.58%	82.03%	84.68%
1.25	90.42%	89.43%	91.6%
1.0	89.31%	89.91%	91.04%

Evaluation 2: Soundness & Completeness

- **Soundness:** Whether a given shallow workflow is meaningful and can help achieve a goal
- **Completeness:** If a given shallow workflow is missing any important actions to achieve a given goal
- Human Subject Evaluation: N = 10

Domain	Soundness	Completeness
Quit Smoking	42%	38%
Couch to 5K	66%	45%
Wedding Planning	36%	43%

Contribution of our Approach

Existing approaches:

- Assume that the model is already given up-front
- Utilize almost completely structured traces to build the model
- Consider a single plan-trace for a given domain

Proposed approach:

- Aggregates experiences of multiple users for a given domain
- Utilizes raw unstructured natural language data
- Extracts structured plan traces by establishing sequentiality between actions
- Builds an approximate causal model
- Robust model that generates meaningful and useful plans

Capturing Preferences: An Extension

- Input
  - Set of plan traces, in natural language, of users performing some activities
  - Social network of people similar and dis-similar to a given person
    - E.g.: Twitter timeline of a person who wants to lose weight
- Output
  - Domain model personalized for a person consisting of state variables (predicates) and actions
- Notes
  - Learned model personalized for an individual, and probabilistically weighted by the choices of others in their social groups (e.g. “near” and “far”)

Prior Work

Prior Art/Criterias	Preferences	Social Network	Learning	No Structured Plan Traces	HTN	AI Planning
HTN Planning (e.g., Nau et al. 2003)	✗	✗	✗	N/A	✔	✔
HTN Planning with Preferences (e.g., Sohrabi et al. 2009)	✔	✗	✗	N/A	✔	✔
Learning HTN or Non-HTN Planning (e.g., Yang et al. 2007)	✗	✗	✔	✗	✔	✔
Learning HTN Planning + Preferences (Li et al., 2012)	✔	✗	✔	✗	✔	✔
Learning Non-HTN Planning + Preferences (e.g., Bryce et al. 2016)	✔	✗	✔	✗	✗	✔
AI Planning with Preferences (e.g., Baier et al., 2009)	✔	✗	✗	N/A	✗	✔
Preference Elicitation/Learning (e.g., Boutilier et al. 2004)	✔	✗	✔	✔	✗	✗
Learning from Social Networks (e.g., Rajaram, et al. 2013)	✔	✔	✔	✔	✗	✗
Proposed Extension of This Work	✔	✔	✔	✔	✔	✔

Conclusions and Future Work

Further Questions & Comments:
Lydia Manikonda
lmanikon@asu.edu

- Extracting models – directly from unstructured data – to support sequential decision making is challenging
- By exploring the challenges and to measure the feasibility of building usable action models, a multi-phase pipeline is proposed
- The pipeline takes unstructured data as input to automatically generate an approximate or incomplete action model
- Evaluations by human raters suggested the plans are sound to a certain extent
- Examining use with current PDDL-style planners future work