

Activity Recognition in TaskTracer and CALO

Tom Dietterich, Victoria Keiser, Jianqiang Shen
School of EECS
Oregon State University

Hung Bui
SRI International



1/25/2011

PAIR 2009



Outline

- TaskTracer: “Project-Oriented User Interface”
 - Semantic Instrumentation
 - Projects & Provenance
- CALO: Activity Recognition
 - Workflow Discovery
 - Activity Recognition

Tasktracer Goals

- Observation: Knowledge workers engage in continuous multi-tasking and interruption recovery
 - Median time between interruptions: 18.0 min
 - Median time to return from interruption: 11.4 min
- **Goal: Support interruption recovery**
- Observation: Knowledge workers manipulate thousands of documents, email messages, and web pages
 - Median 3,909 over 4 weeks
- **Goal: Support finding and re-finding relevant information**
 - On average, each item is opened 1.74 times over 4 weeks

TaskTracer Hypotheses

- **Hypotheses:**
 - Users find it natural to organize their work in terms of projects
 - Each document, web page, email message, and person is associated with a small number of projects
 - User time at the desktop can be viewed as multi-tasking among a set of active projects
 - When working on project P, users will tend to access only other items also associated with P
- **Implications:**
 - Ask users to define a hierarchy of projects
 - Track the user's current project P
 - Automatically tag new items based on the user's current project P
 - Support information access to existing items based on the current project P

Semantic Instrumentation

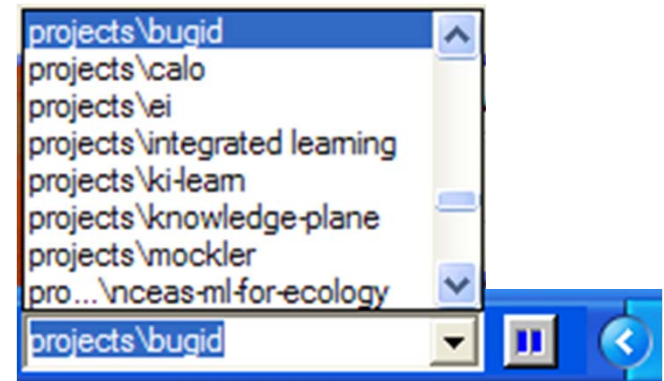
- **Applications:**
 - Word, Excel, PowerPoint, Outlook, Internet Explorer, Windows Explorer, GSView, Acrobat, Visual Studio, Thunderbird
- **Application Events:**
 - Documents: New, Change, Open, Save, Save As, Close
 - Email: Open, Close, Send, Reply, Forward, Attach File, Save Attachment, Open Attachment, Incoming Email, Click on Hyperlink
 - Web pages: Open, Navigate, Upload File, Download File
- **OS Events:**
 - File Create/Delete/Rename
 - Window Focus
 - Copy/Paste
 - Suspend/Resume/Idle

Tracking the User's Current Project

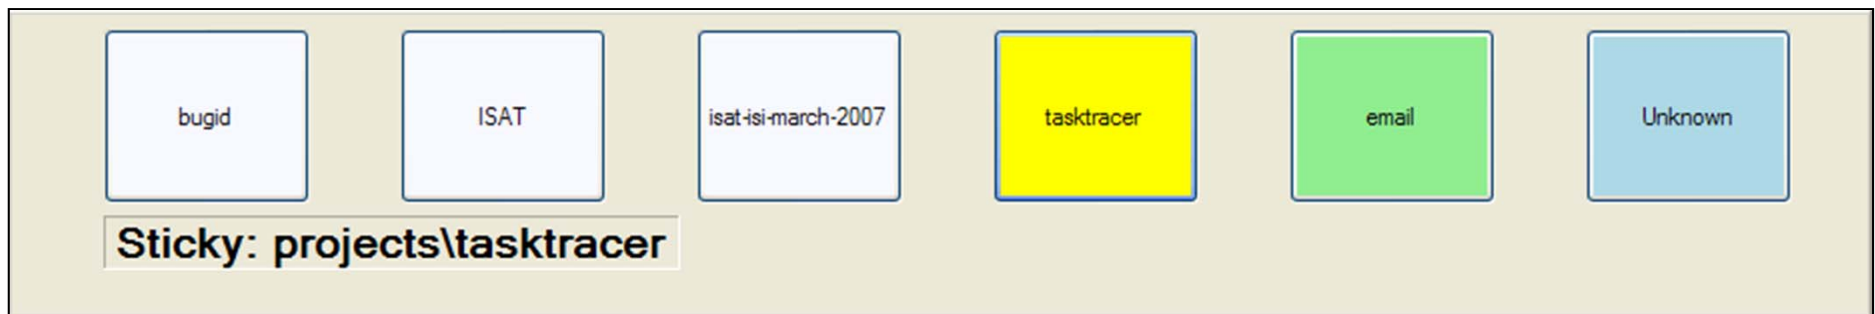
- Manual input from user: Project Selector
- Incoming email classifier: Email Predictor
- Project Switch detector: Project Predictor

Manual Input from User

- Project Selector: Windows Task Bar:



- Pop-up menu of recent projects: Control+tilde



Email Predictor

- Predicts most likely project of each email message

Predicted project is added as a "Category" in Outlook

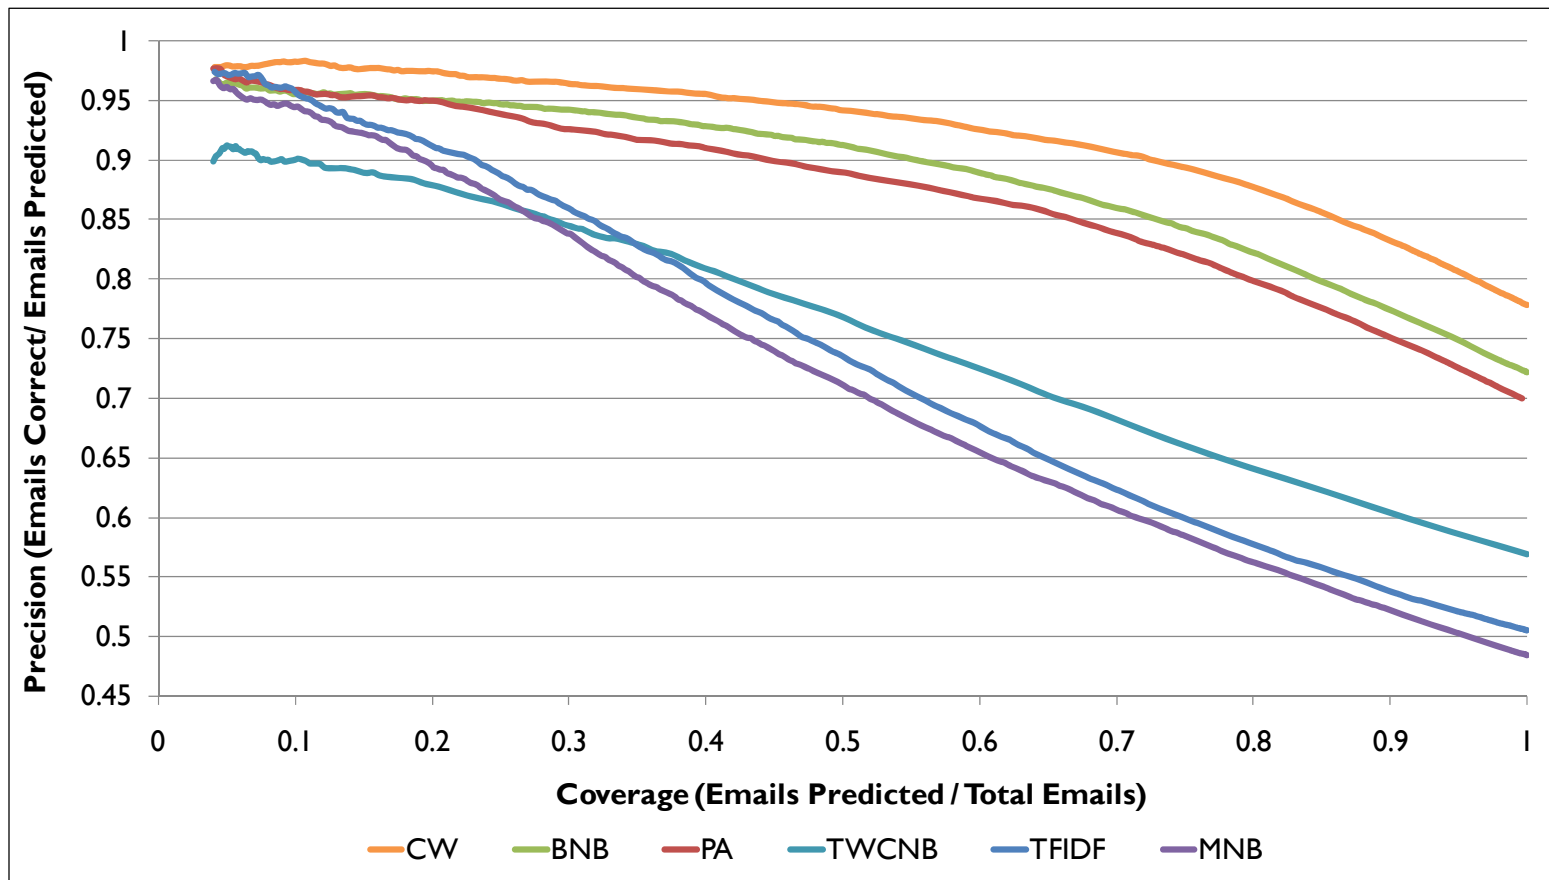
The screenshot shows an Outlook inbox with a table of email messages. The table has columns for From, Subject, Categories, and Received. The messages are grouped by date: Sunday, Yesterday, and Today. A red callout box points to the 'Categories' column, stating that the predicted project is added as a 'Category' in Outlook.

From	Subject	Categories	Received
Sandra Zilles	Re: Your United flight confirmation - June 18, 2008 - Portlan...	conf\alberta-june-2008	Sat 4/19/200...
Julia Jones	References for igert renewal pre-proposal	grant\ysf\igert-2008	Sat 4/19/200...
Date: Sunday			
Tracy Kugler	CS student collaborator and literature for amenities-develop...	students\kugler	Sun 4/20/20...
Julia Jones	Draft of pre-proposal for your review	grant\ysf\igert-2008	Sun 4/20/20...
Karen Utgoff	Directions to our house	friends\utgoff	Sun 4/20/20...
Julia Jones	Ei renewal pre-proposal draft reference list	grant\ysf\igert-2008	Sun 4/20/20...
Date: Yesterday			
technews	ACM TechNews; Monday, April 21, 2008	reading	Mon 4/21/20...
Martin Hofmann	[gila-sub-pi] GILA slides	grant\darpa\integrated-l...	Mon 4/21/20...
Xinlong Bao	Results from the latest TTPCE run	projects\calo\jr	Mon 4/21/20...
Mike Wynkoop	Re: no progress on paper today	projects\tl	Mon 4/21/20...
Carlos Jensen	Re: FW: MURI Human Subjects	grant\muri-2007	Mon 4/21/20...
Aron Culotta	Re: Dinner tomorrow	conf\umass-april-2008	Mon 4/21/20...
Carol Rivin	Re: visit?	family\carol	Mon 4/21/20...
Edward Waymire	Re: Draft of pre-proposal for your review	grant\ysf\igert-2008	Mon 4/21/20...
Xinlong Bao	Re: Counting Weight Question	projects\calo\jr	Mon 4/21/20...
Date: Today			
Thomas Hofmann	Review Request, JAIR Submission 2536	service\JAI	Tue 4/22/20...
Margaret Burnett	Re: MURI Human Subjects	grant\muri-2007	Tue 4/22/20...
Ahmed Abdelhakim	Re: SPT 102, 103 & 104	comp\intellect desk t	Tue 4/22/20...

Email Predictor: Confidence-Weighted Classifier

- Drezde, Crammer & Pereira, ICML 2008
 - Linear classifier
 - Maintains diagonal Gaussian distribution over the weight vector
 - When a mistake is made
 - Update the Gaussian so that with probability $1 - \alpha$, a randomly sampled weight vector would not have made the mistake
- Features:
 - sender
 - set of all recipients
 - words in Subject:
 - words in body

Email Classifier Performance



- Dietterich's email 2004 – 2008
- almost 21,000 examples
- 381 classes ranging in size from 1 to more than 2500 messages


Predicting Project Switches

- Extract the following information once per minute:
 - Current active “resource”
 - Window title
 - Pathname / URL
 - Current declared project
 - Titles of all windows
 - Resources in all windows

Project Predictor Features

- Project-Specific Features $F_p(y)$
 - Strength of association of active resource with y
 - % of open resources associated with y
 - Importance of title word x to project y
 - Modified TF-IDF score
- Switch-Specific Features $F_s(y, y')$
 - # of resources closed in last minute
 - % of open resources accessed in last minute
 - Time since user's last explicit switch

Project Predictor



- Let
 - $F_p(X,y)$ = project-specific features
 - $F_S(X_t, X_{t+1})$ = switch-specific features
 - Λ = learned parameters

- Scoring function:

$$G(y_t, y_{t+1} | X_t, X_{t+1}) = \Lambda_1 \cdot F_P(X_t, y_t) + \Lambda_2 \cdot F_P(X_{t+1}, y_{t+1}) + I(y_t \neq y_{t+1}) [\Lambda_3 \cdot F_S(X_t, X_{t+1})]$$

Regularized Passive-Aggressive Algorithm

(Crammer et al., 2006; Shen, 2008)

- Add a penalty on the size of Λ
- Closed-form solution:

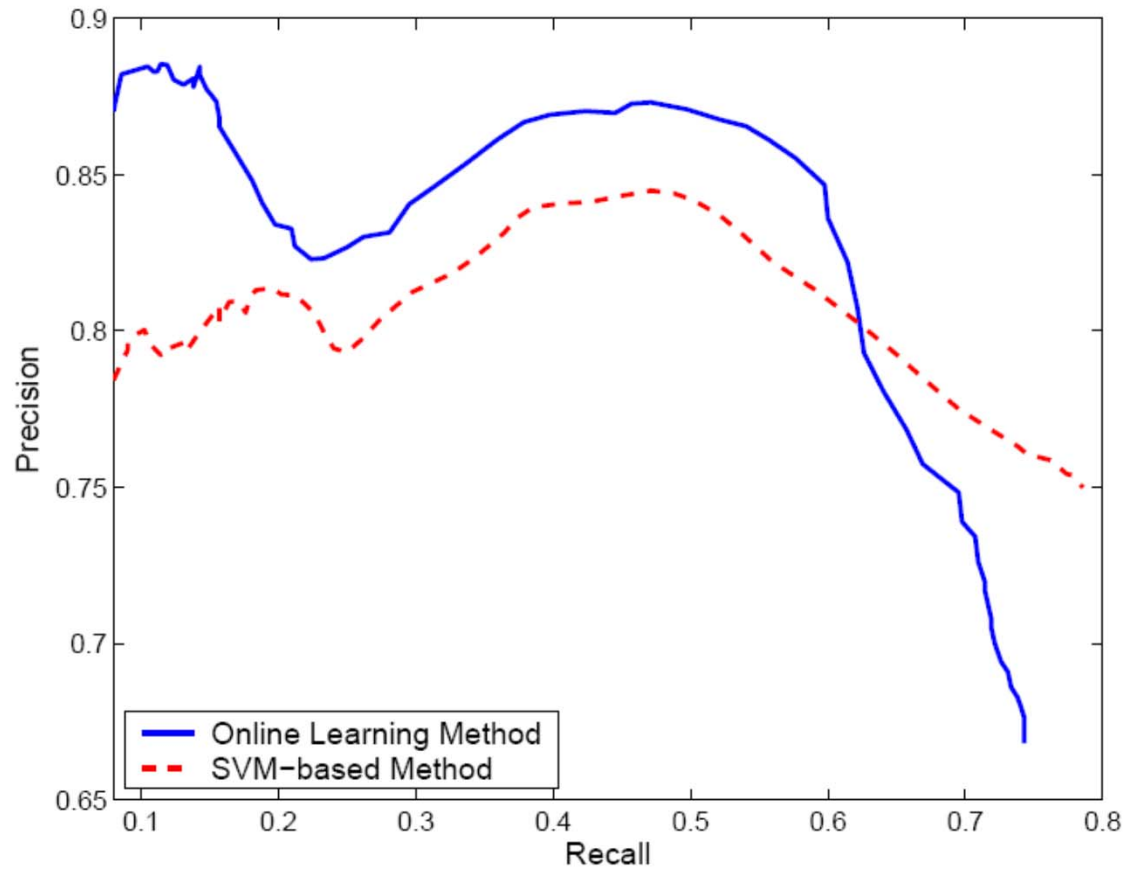
$$\Lambda_{t+1} = \arg \min_{\Lambda \in \mathbb{R}^n} \frac{1}{2} \|\Lambda - \Lambda_t\|_2^2 + C\xi^2 + \frac{\alpha}{2} \|\Lambda\|_2^2$$

subject to $g(\langle y_1, y_2 \rangle) - g(\langle \hat{y}_1, \hat{y}_2 \rangle) \geq 1 - \xi$.

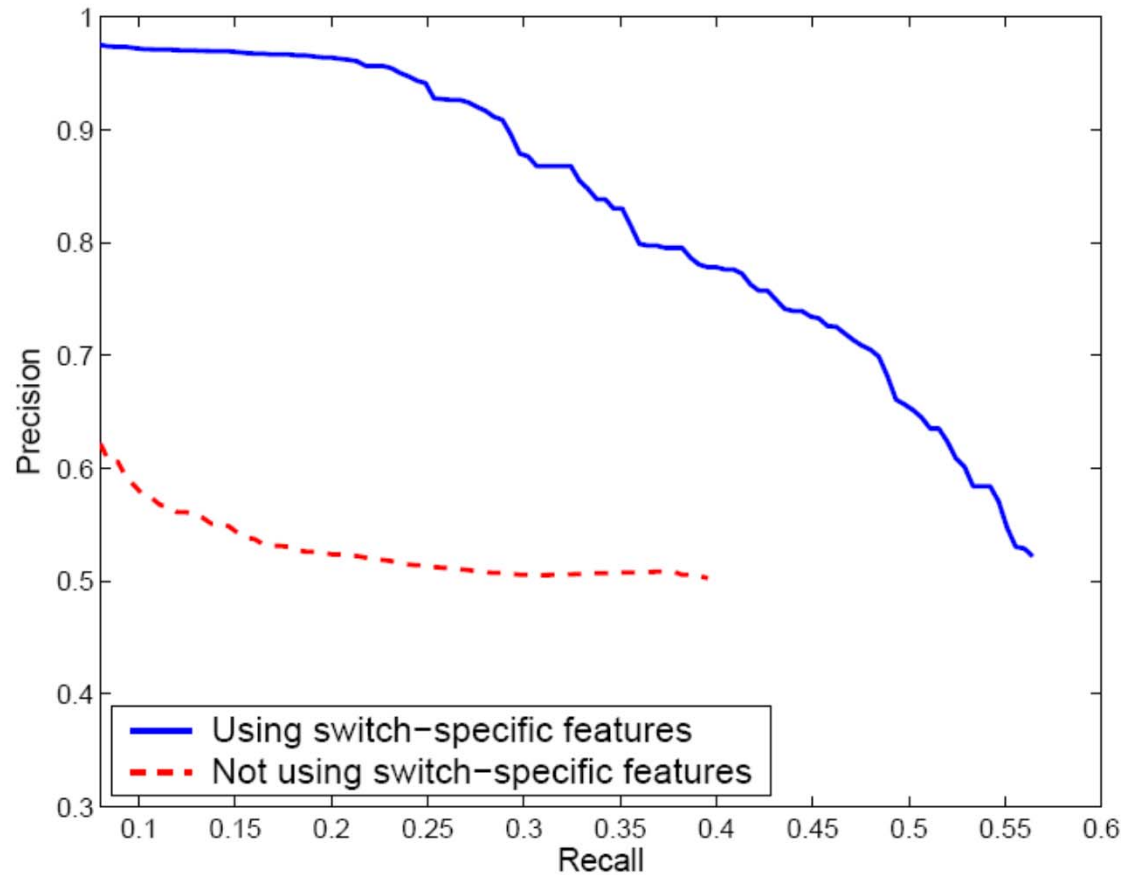
$$\Lambda_{t+1} := \frac{1}{1 + \alpha} (\Lambda_t + \tau_t \mathbf{Z}_t),$$

$$\tau_t = \frac{1 - \Lambda_t \cdot \mathbf{Z}_t + \alpha}{\|\mathbf{Z}_t\|_2^2 + \frac{1+\alpha}{2C}}.$$

Results on My Data



Results on Another User



Project-Based Services

- Task explorer
- Folder predictor
- Outlook search folders
- Time tracking
- Project-specific notes

Task Explorer

Supports Interruption Recovery and Information

All folders, documents, web pages, emails, and contacts

Project Hierarchy

The screenshot shows the Task Explorer application window titled "Task Explorer - projects\tasktracer". The window has a menu bar with "View", "Tools", and "Help". Below the menu bar is a toolbar with icons for "All", "Web", "Folders", "Word", "Excel", "PowerPoint", "Mail", "Contacts", and "Other".

The left pane displays a project hierarchy with a tree view. The "projects" folder is expanded, showing sub-items like "bugid", "calo", "ei", "integrated learning", "ki-learn", "ssl", "tasktracer", and "tdas-study".

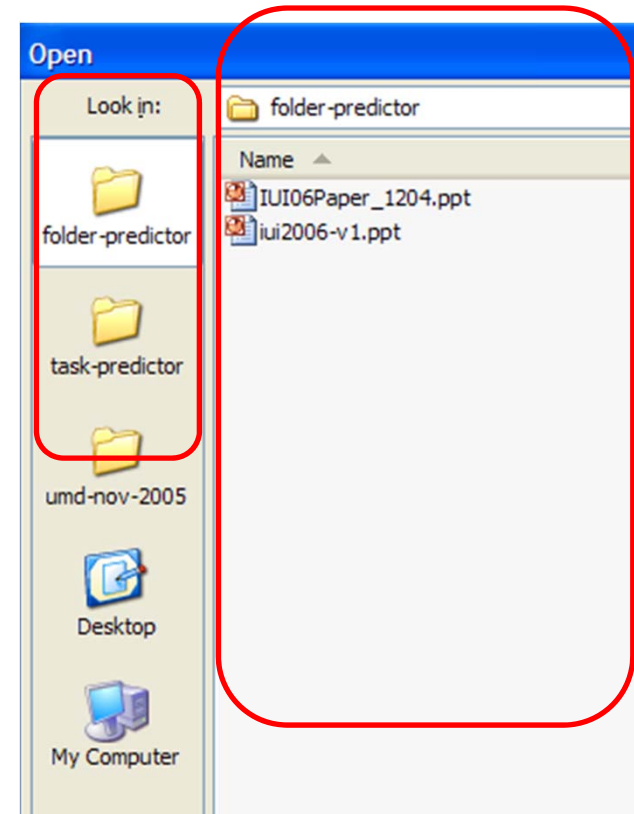
The right pane displays a list of files and folders. The list has columns for "Name", "Path", "Type", and "Date Modified".

Name	Path	Type	Date Modified
talks	Z:\p\tasktracer\talks\	Folder	9/19/2006
tasktracer-rochester-v1.ppt	Z:\...ktracer\talks\tasktracer-rochester-v1.ppt	PowerPoint	9/19/2006
figure	Z:\p\tasktracer\papers\ijcai07\v3\figure\	Folder	9/19/2006
user_pvc.eps	z:\...er\papers\ijcai07\v3\figure\user_pvc.e...	File	9/19/2006
user_coap.eps	z:\...r\papers\ijcai07\v3\figure\user_coap.e...	File	9/19/2006
art_coap.eps	z:\...er\papers\ijcai07\v3\figure\art_coap.eps	File	9/19/2006
coverage.eps	z:\...er\papers\ijcai07\v3\figure\coverage....	File	9/19/2006
v3	Z:\p\tasktracer\papers\ijcai07\v3\	Folder	9/19/2006
ijcai07	Z:\p\tasktracer\papers\ijcai07\	Folder	9/19/2006
papers	Z:\p\tasktracer\papers\	Folder	9/19/2006
tasktracer	Z:\p\tasktracer\	Folder	9/19/2006
tdas-final-report-v1.ppt	Z:\p\tdas\tdas-final-report-v1.ppt	PowerPoint	9/19/2006
tasktracer-gatech-v1.ppt	Z:\p\tasktracer\talks\tasktracer-gatech-v1....	PowerPoint	9/19/2006
talks	Z:\p\bugs\talks\	Folder	9/19/2006
usable-ml	Z:\p\tasktracer\papers\iui2007\usable-ml\	Folder	9/17/2006
IUI07-060917-1521.doc	Z:\...2007\usable-ml\IUI07-060917-1521.doc	Word	9/17/2006
INTEL-09-15-2006	C:\...ded-files\tasktracer\INTEL-09-15-2006\	Folder	9/16/2006
tasktracer	C:\downloaded-files\tasktracer\	Folder	9/16/2006
IUI07-060915_1400.doc	Z:\...2007\usable-ml\IUI07-060915_1400.doc	Word	9/15/2006
IUI07-060915_0957.doc	Z:\...2007\usable-ml\IUI07-060915_0957.doc	Word	9/15/2006
http://www.cs.orst.edu/~tgd/	http://www.cs.orst.edu/~tgd/	IE Explorer	9/14/2006
http://citeseer.ist.psu.edu/	http://citeseer.ist.psu.edu/	IE Explorer	9/14/2006
cannot find server	http://citeseer.ist.psu.edu/segal99mailcat.ht...	IE Explorer	9/14/2006

The status bar at the bottom shows the date "1/25/2011" and the page number "18".


Information Re-Finding: Folder Predictor: Shortcuts to Relevant Folders

- Maintain statistics on file opens and saves on a per-project basis
 - Recency-weighted count of saves and opens
- When user initiates open/save compute 3 folders to minimize expected number of clicks to get to the desired folder



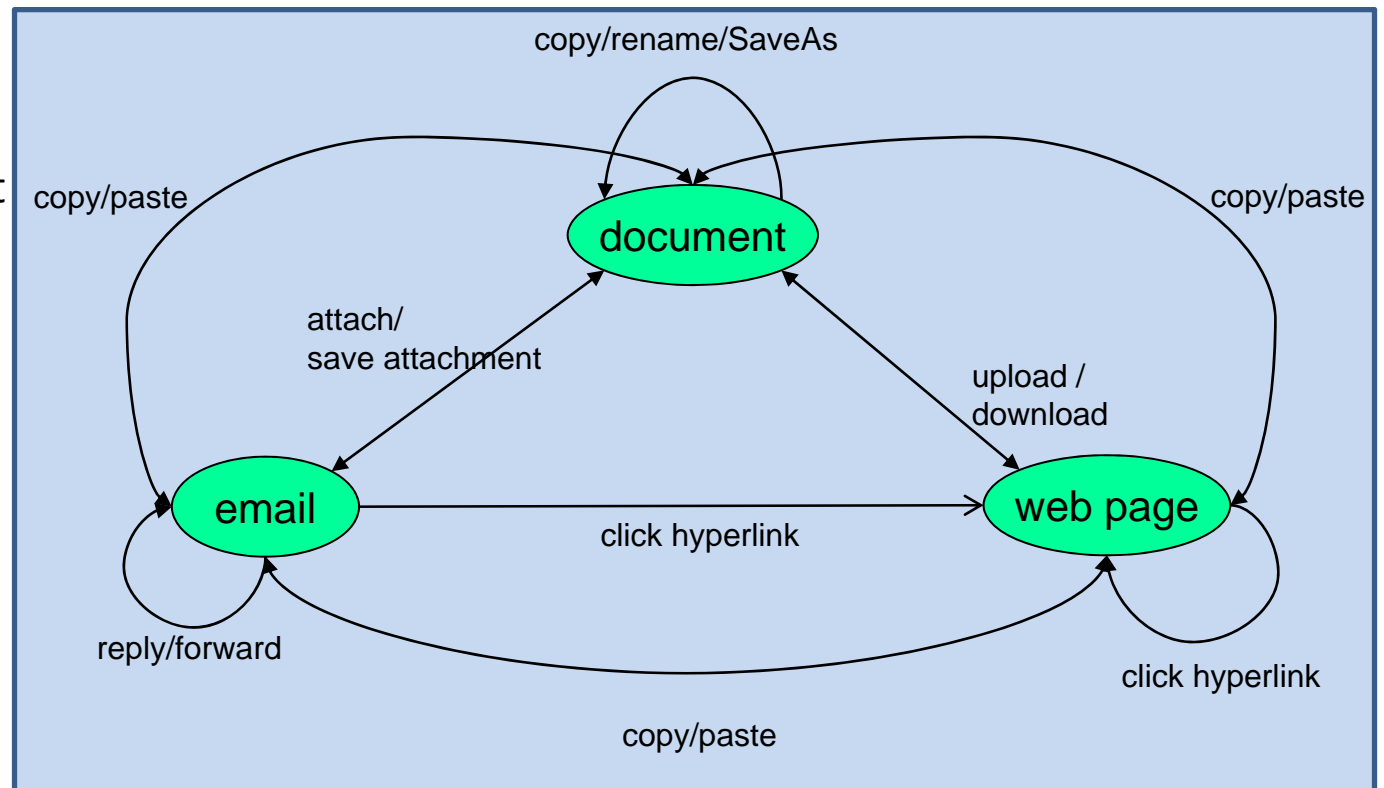
$$\operatorname{argmin}_{\{f_1, f_2, f_3\}} \sum_f P(f \mid \text{proj}) \cdot \min \{ \text{clicks}(f_1, f), 1 + \text{clicks}(f_2, f), 1 + \text{clicks}(f_3, f) \}$$

Average Cost to Reach Target Folder



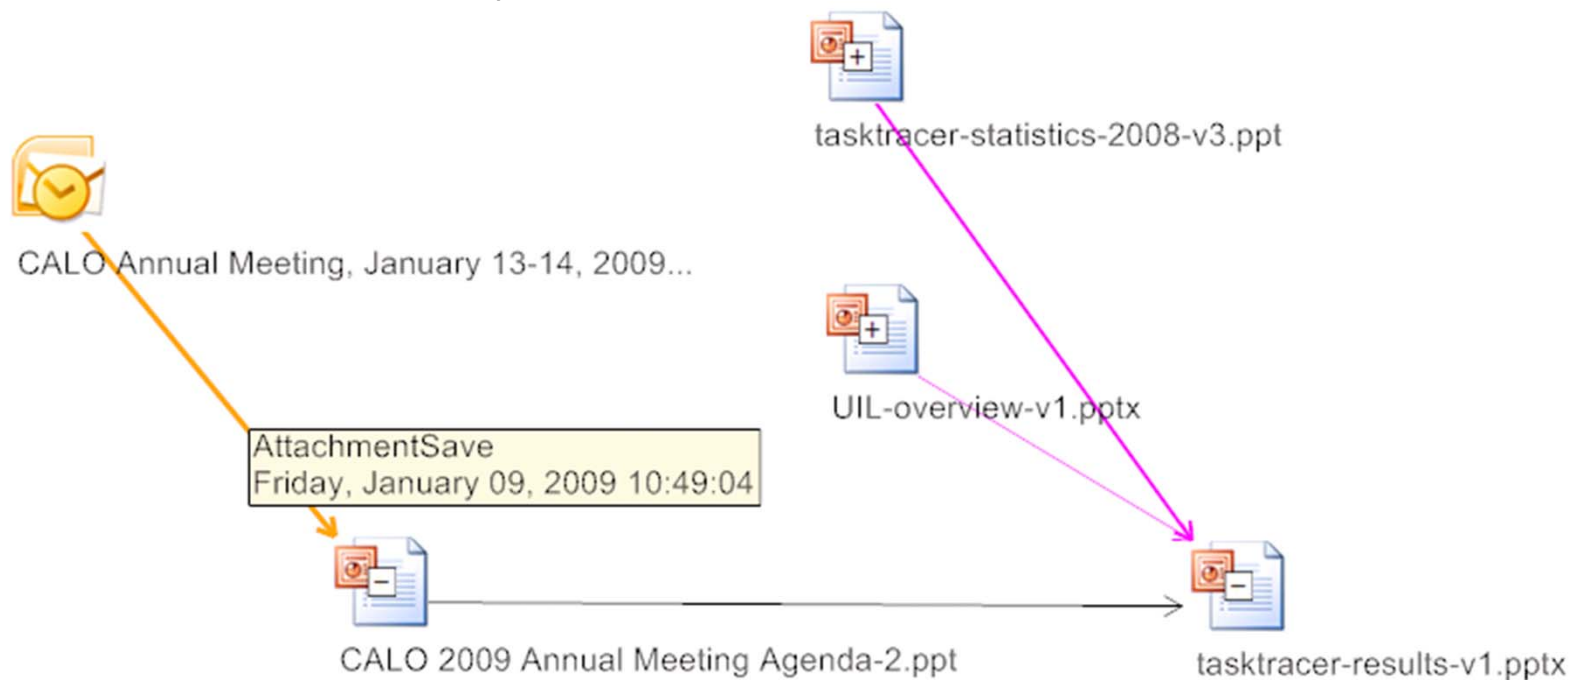
Provenance Instrumentation

- File-to-File
 - copy/paste
 - cut/paste
 - SaveAs
 - file system copy
 - file system rename
- Email-to-File
 - save attachment
 - add attachment
- Email-to-Email
 - reply
 - forward
- Web-to-file
 - download file
 - upload file
- Email-to-Web
 - click on hyperlink
- Web-to-Web
 - click on hyperlink



Provenance-Based Information Access

- Right-click on object opens Provenance Graph
 - email header in Outlook
 - attachment in Outlook
 - file name in Windows Explorer



Intel Smart Desktop Study


- 6 participants (~4 weeks each)
- Beta 3 of SmartDesktop (with additional provenance instrumentation)

Evaluating the TaskTracer Hypotheses

- Hypothesis: When working on project P, users will tend to access only other items also associated with P
- Identify all cases where user switches between two items A and B and classify them:

	A→B copy/paste	B changed	None
A & B in same project	Within-Project Information Access		Ignore
A & B in different projects	Cross-Project Information Access	Weak Cross-Project Information Access	Multi-tasking switch

Cross-Project Information Access is Rare



TaskTracer Summary

- “Project”: very high level and very useful notion of “activity”
 - **Interruption recovery**
 - **Information re-finding**
- Mix of user input and automated activity recognition
- Online classifiers that are responsive to user feedback
- Integration with Windows XP and Office
- Open Source release this fall

Outline

- TaskTracer: “Project-Oriented User Interface”
 - Semantic Instrumentation
 - Projects & Provenance
- **CALO: Activity Recognition**
 - Workflow Discovery
 - Activity Recognition

Discovering Workflows

- Build an information flow graph
 - Node: each visited object (files, emails, webpages)
 - Arcs: TaskTracer provenance links + other inferred connections
 - Collapse certain patterns
- Search for the frequent subgraphs in the information flow graph
 - Step 1: Ignore edge labels
 - Step 2: Dynamic program to find most frequent labeled subgraphs
- Some assumptions
 - A workflow is a (weakly) connected graph
 - A workflow involves at least $k = 3$ resources
 - A workflow involves at least one non-email resource

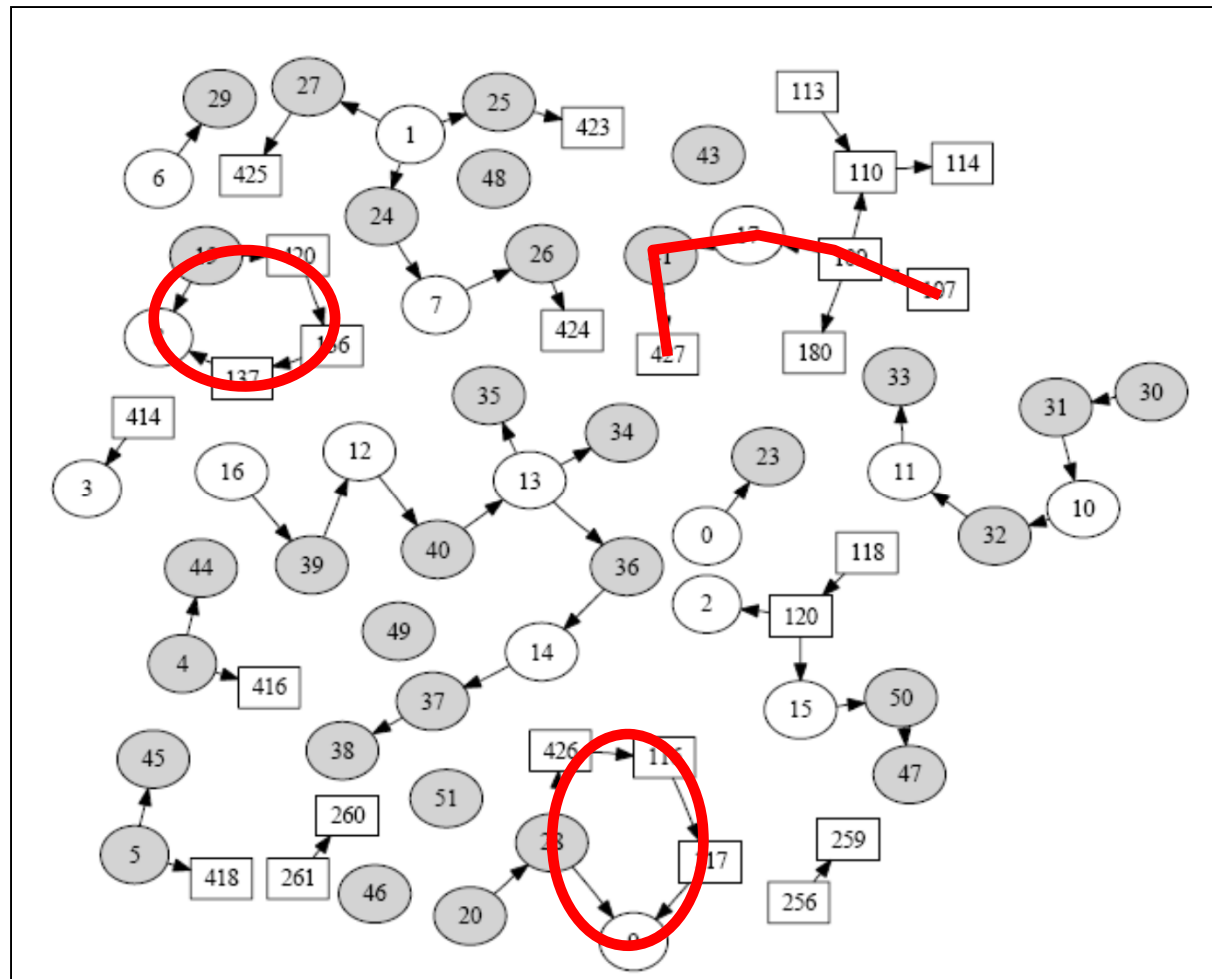
Inferred Relationships

- Document conversion
 - convert to PDF file
 - zip and unzip
- Email reference
 - email body refers to a document (e.g., by title)

Collapse Certain Patterns

- File editing
 - Replace chain of SaveAs links with SaveAs*
- Email conversation
 - Replace Send/Receive/ReplyTo links with ReplyTo*

Resulting Information Flow Graph




Mining Frequent Subgraphs

- Construct frequent pattern candidates from the existing ones [Nijssen & Kok: KDD04]
- Restrict to “Closed” patterns: Can’t add a node or edge without shrinking the set of covered instances

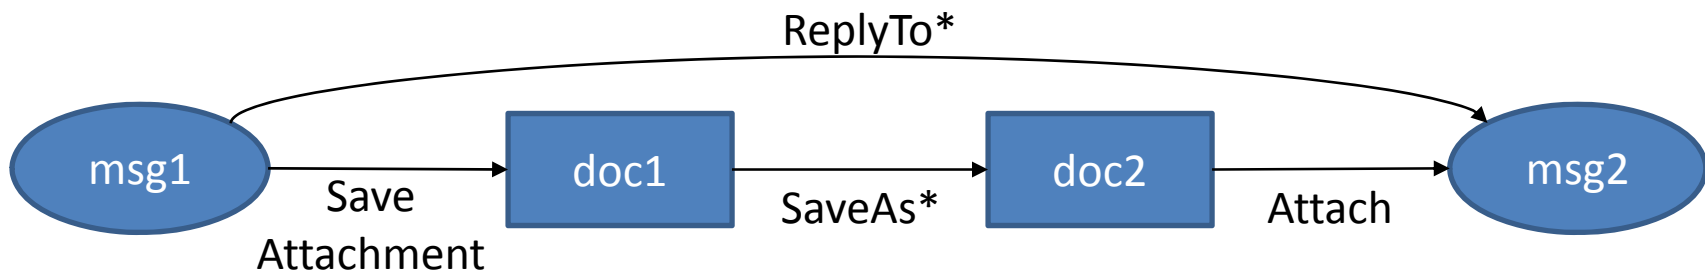
Assigning Action Type Labels to Each Edge

- Searching For Frequent Action Paths
 - Choose action types for each edge in the discovered subgraphs
 - Simply selecting the most frequent action for each edge does not work



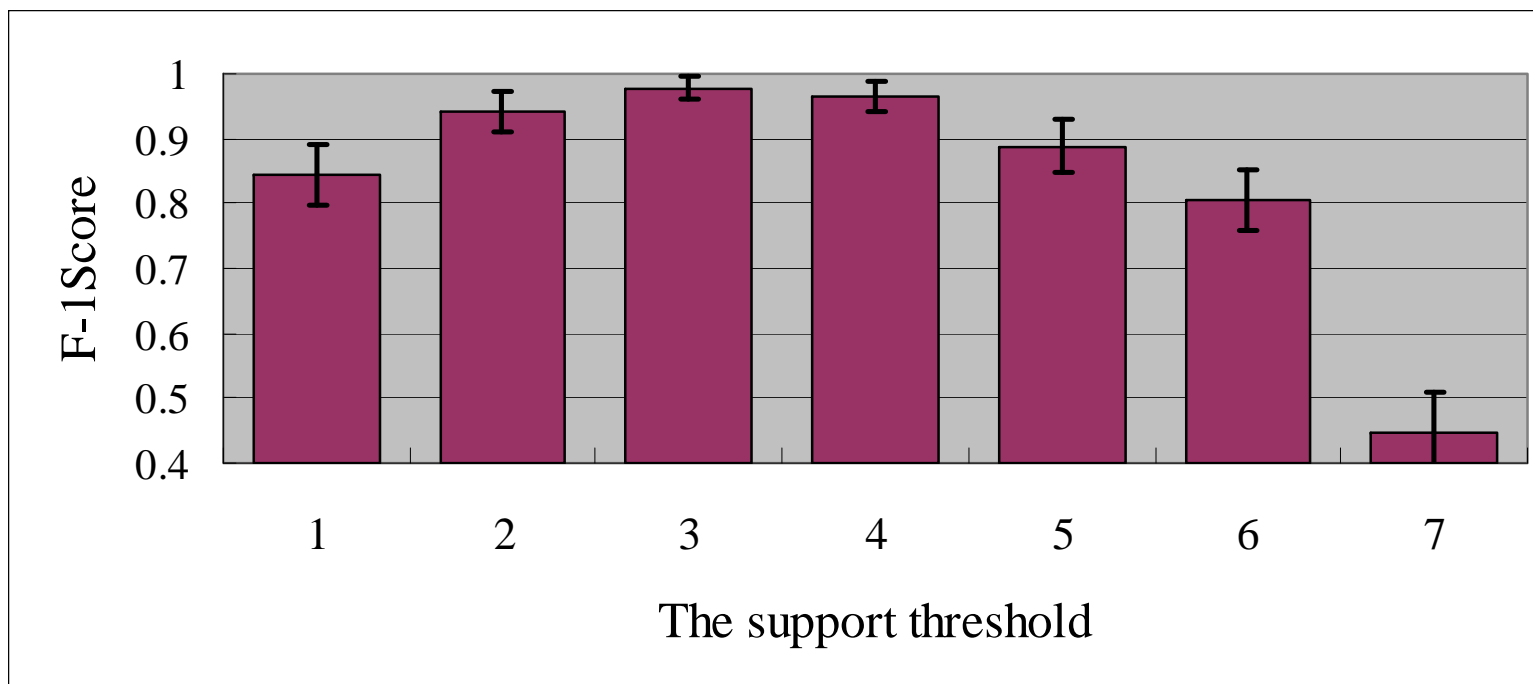
- Efficient dynamic programming algorithm

Example Discovered Workflow



- CALO experiment
 - four participants
 - preparing and reviewing papers for a conference
 - filing travel requisitions and reimbursements
 - web browsing, etc.
- Five workflows discovered

Results on CALO Data



Recall: the number of correctly discovered resources and actions divided by the number of real resources and actions

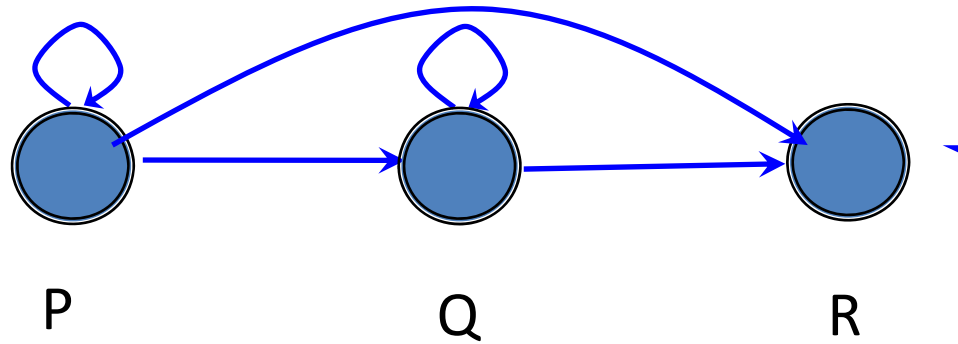
Precision: the number of correctly discovered resources and actions divided by the number of predicted resources and actions

F1: $2 * \text{precision} * \text{recall} / (\text{precision} + \text{recall})$


Real-Time Recognition of Workflows

- **Problem:** recognize current user's workflow(s), track their progress and their parameters
- **Challenge:**
 - Workflow instances vary (optional steps, non-deterministic branching and ordering)
 - Workflows have parameters (e.g. filename)
 - Huge number of objects (emails, files, contacts)
 - Workflow steps need to be separated from background activities
 - Multiple workflows need to be separated from each other

Tracking a Single Workflow: Logical Hidden Markov Model




Logical Hidden Markov Model



- Transition probabilities represented as logical probabilistic rules

$$\exists x; y; z \ P(x; y) \wedge Q(y; z); w = 0:5=jzj$$

Logical Hidden Markov Model



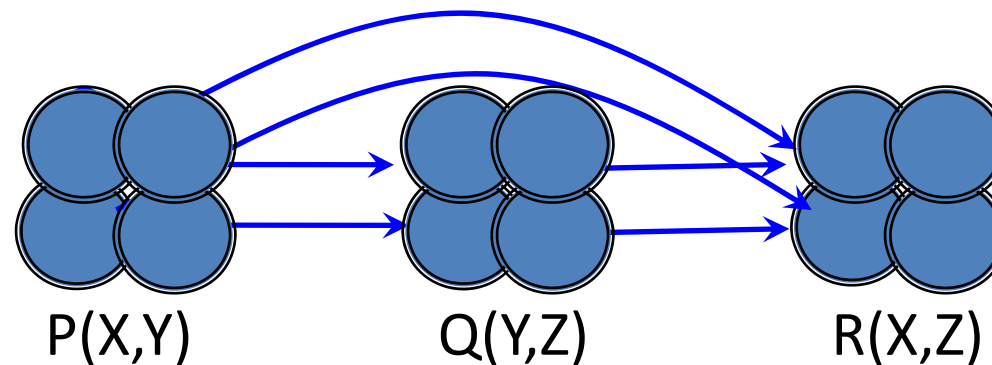
- Transition probabilities represented as logical probabilistic rules

$$\exists x; y; z \ P(x; y) \wedge Q(y; z); w = 0.5 = |z|$$

Ground HMM

$P(a; a) \wedge Q(a; c); w = 0.5 = z $
$P(a; b) \wedge Q(b; c); w = 0.5 = z $
$P(b; a) \wedge Q(a; c); w = 0.5 = z $
$P(b; b) \wedge Q(b; c); w = 0.5 = z $
\vdots

Inference in Logical HMMs



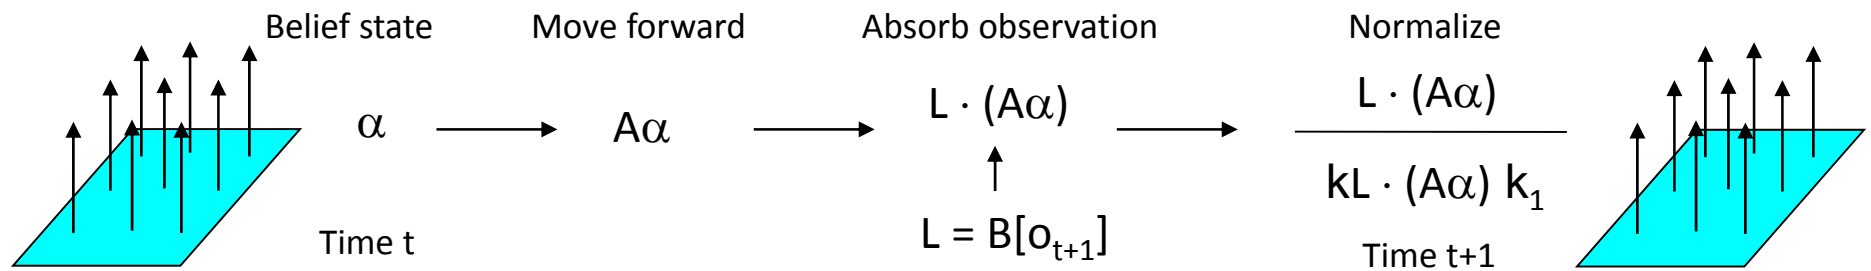
- **Naïve method**

- Convert to ground HMM
- Complexity $O(S^2)$
 - $S = (\# \text{ objects})^{\text{arity}}$
- 1000 objects, arity=3
 - $S^2 = 10^{18}$
- Hopeless!

- **Kersting et al 2006**

- Incremental construction of ground HMM
- Only generate the set of ground states feasible/reachable given the starting state and the observations
- Problematic when too many feasible states remain

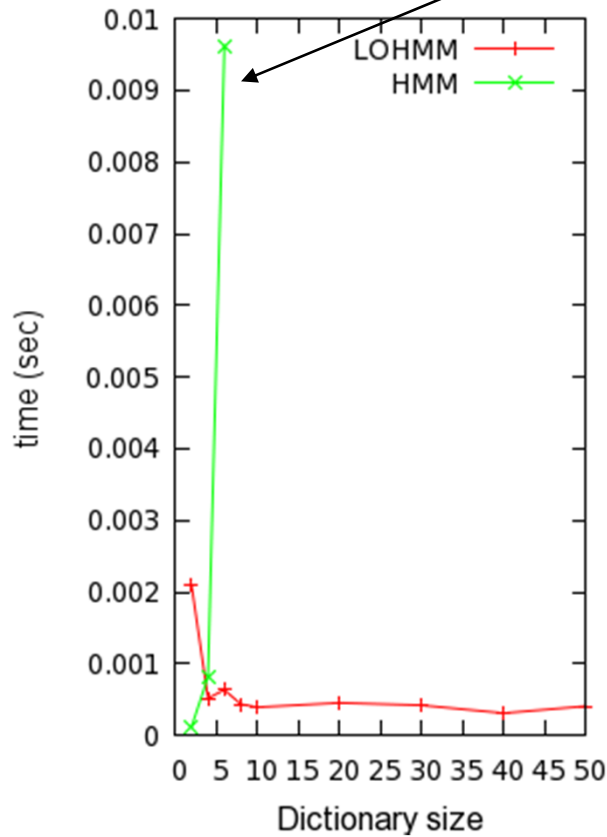
Lifted Inference for Logical HMMs




- Three main operators
 - Matrix multiplication
 - Hadamard product
 - Normalization (L1-norm)
- These can be implemented for LHMMs without grounding

Results on Randomly Generated Models

Ground inference via HMM is hopeless

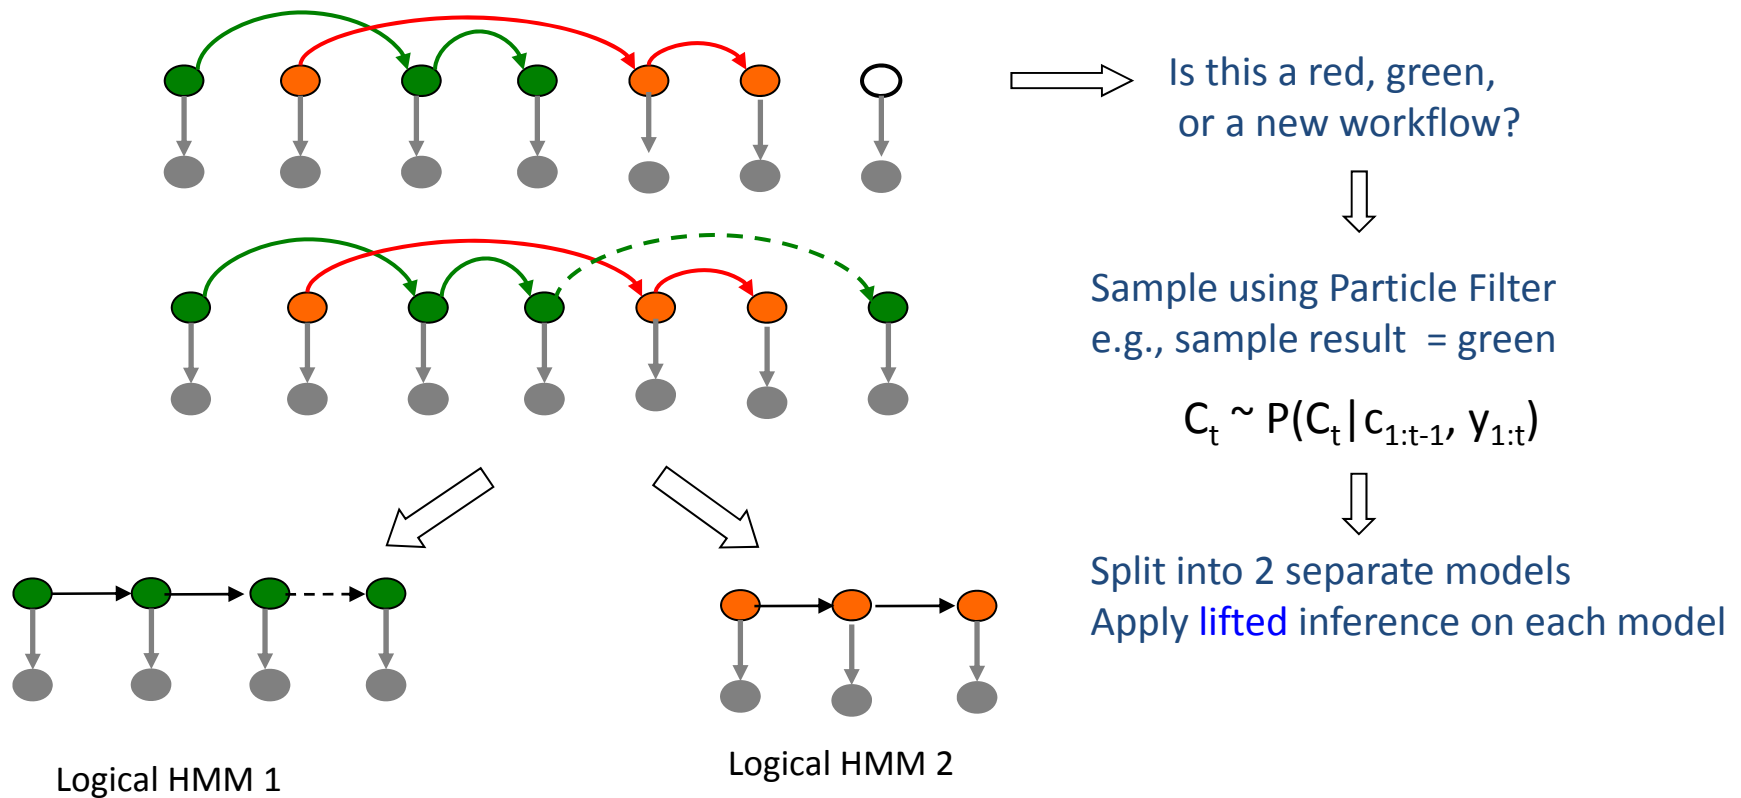


Lifted inference almost independent of domain size



Lifted inference is linear in # variables in predicate

Tracking Multiple, Interleaved Workflows: Lifted Rao-Blackwellized Particle Filter



- Sample only one variable (which workflow is active)
 - All the other random variables (workflow states, workflow parameters) are marginalized out via lifted inference

Experimental Test with Learned Workflows

- Leave one user out
- At each time t ,
 - WARP makes prediction
 - We score prediction for correctness
 - We provide corrective feedback
- Results
 - Precision: 91.3%
 - Recall: 66.7%
 - F1: 77.1%
- Major bug: poor accuracy when recognizing the first step in a new workflow

Open Problems

- **Blackbox Semantic Instrumentation**
 - Can we infer semantic events from easily-observed events (system calls, file system operations, window manager operations)?
- **Deeper Semantic Analysis of Email and Documents**
 - Recognize speech acts, discussions, decisions
- **Detect Start of New Workflow**
 - May require multiple events to detect (“smoothing” instead of “filtering”)
- **Usable Integration with TaskTracer**

Acknowledgements

- Funding:
 - NSF MKIDS program grant IIS-0133994
 - DARPA PAL program
 - This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. FA8750-07-D-0185/0004. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the DARPA, or the Air Force Research Laboratory (AFRL).
 - Intel R&D Council