Discovering Latent Structure in Task-Oriented Dialogues

Ke Zhai and Jason D. Williams

June 22, 2014

Dept. of Computer Science
University of Maryland
College Park
zhaike@cs.umd.edu

Microsoft Research
Redmond
Jason.Williams@microsoft.com

[1] Internship work at Microsoft research, Redmond

Dialogue Systems and State Tracking

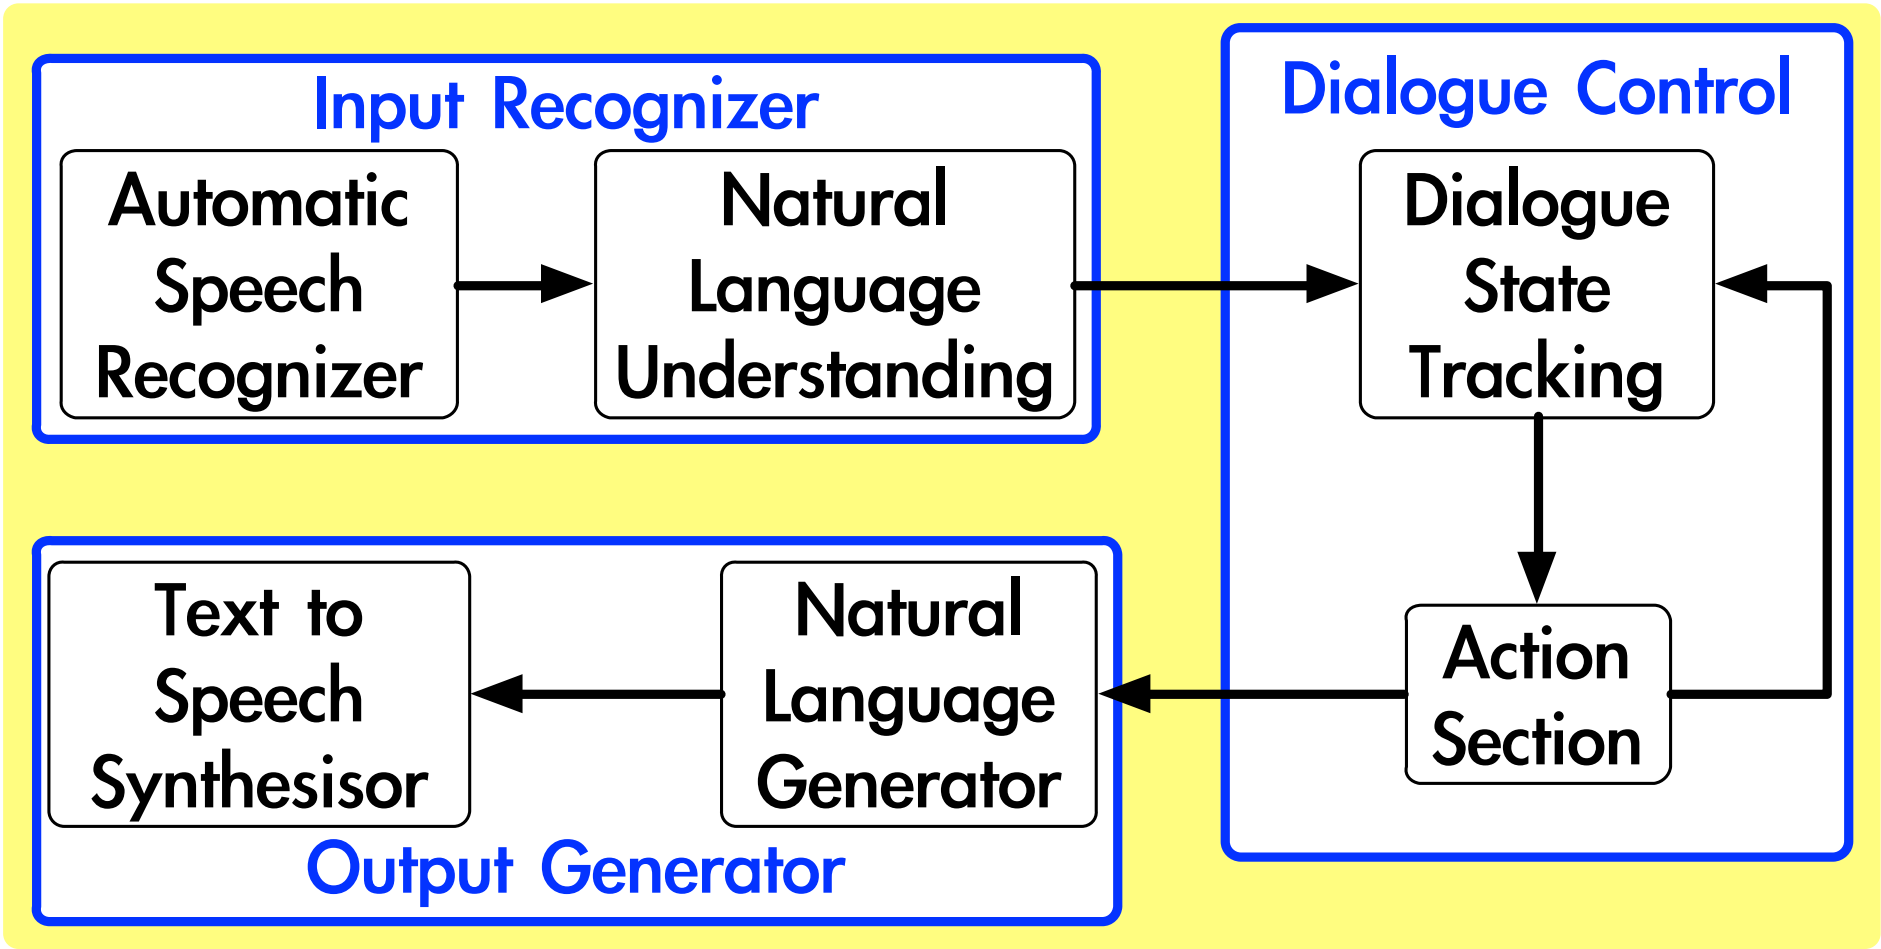

- Input Recognizer: speech-to-text and sematic analysis
- Dialogue Control: flow control and decision making
- Output Generator: natural language and speech generator

Dialogue Systems and State Tracking

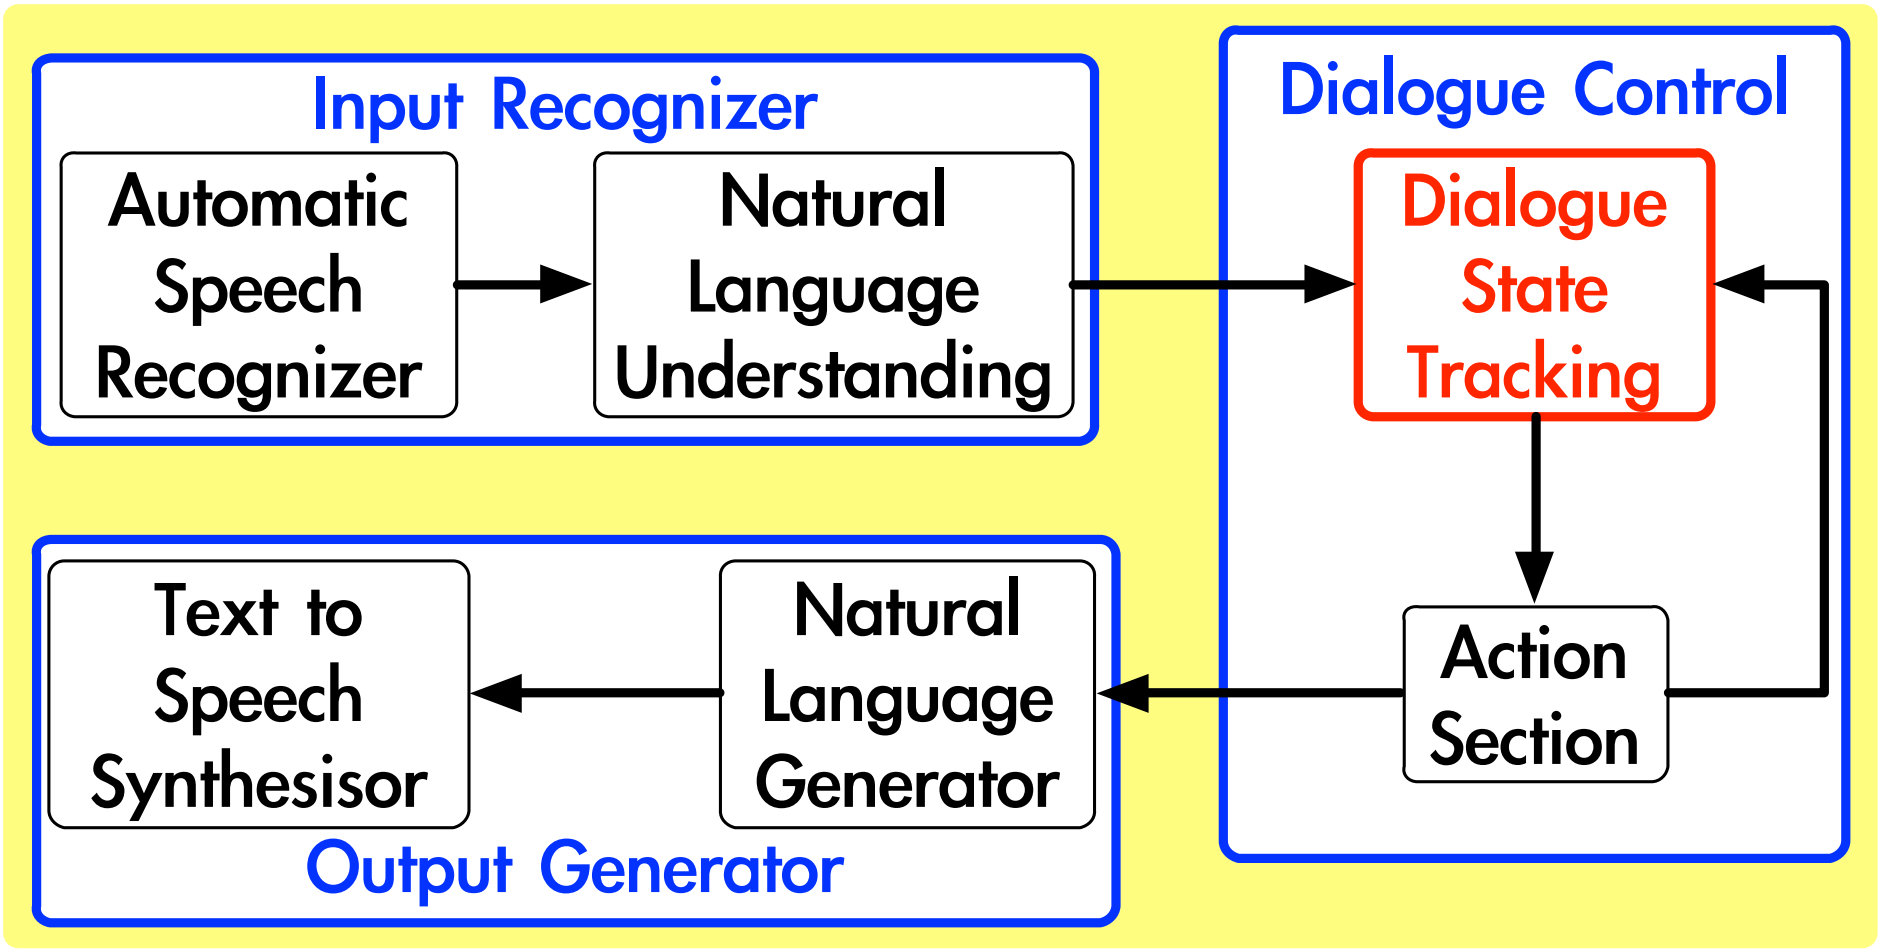

Key Challenge

tracking dialogue states and understanding conversation flows

Dialogue Systems and State Tracking

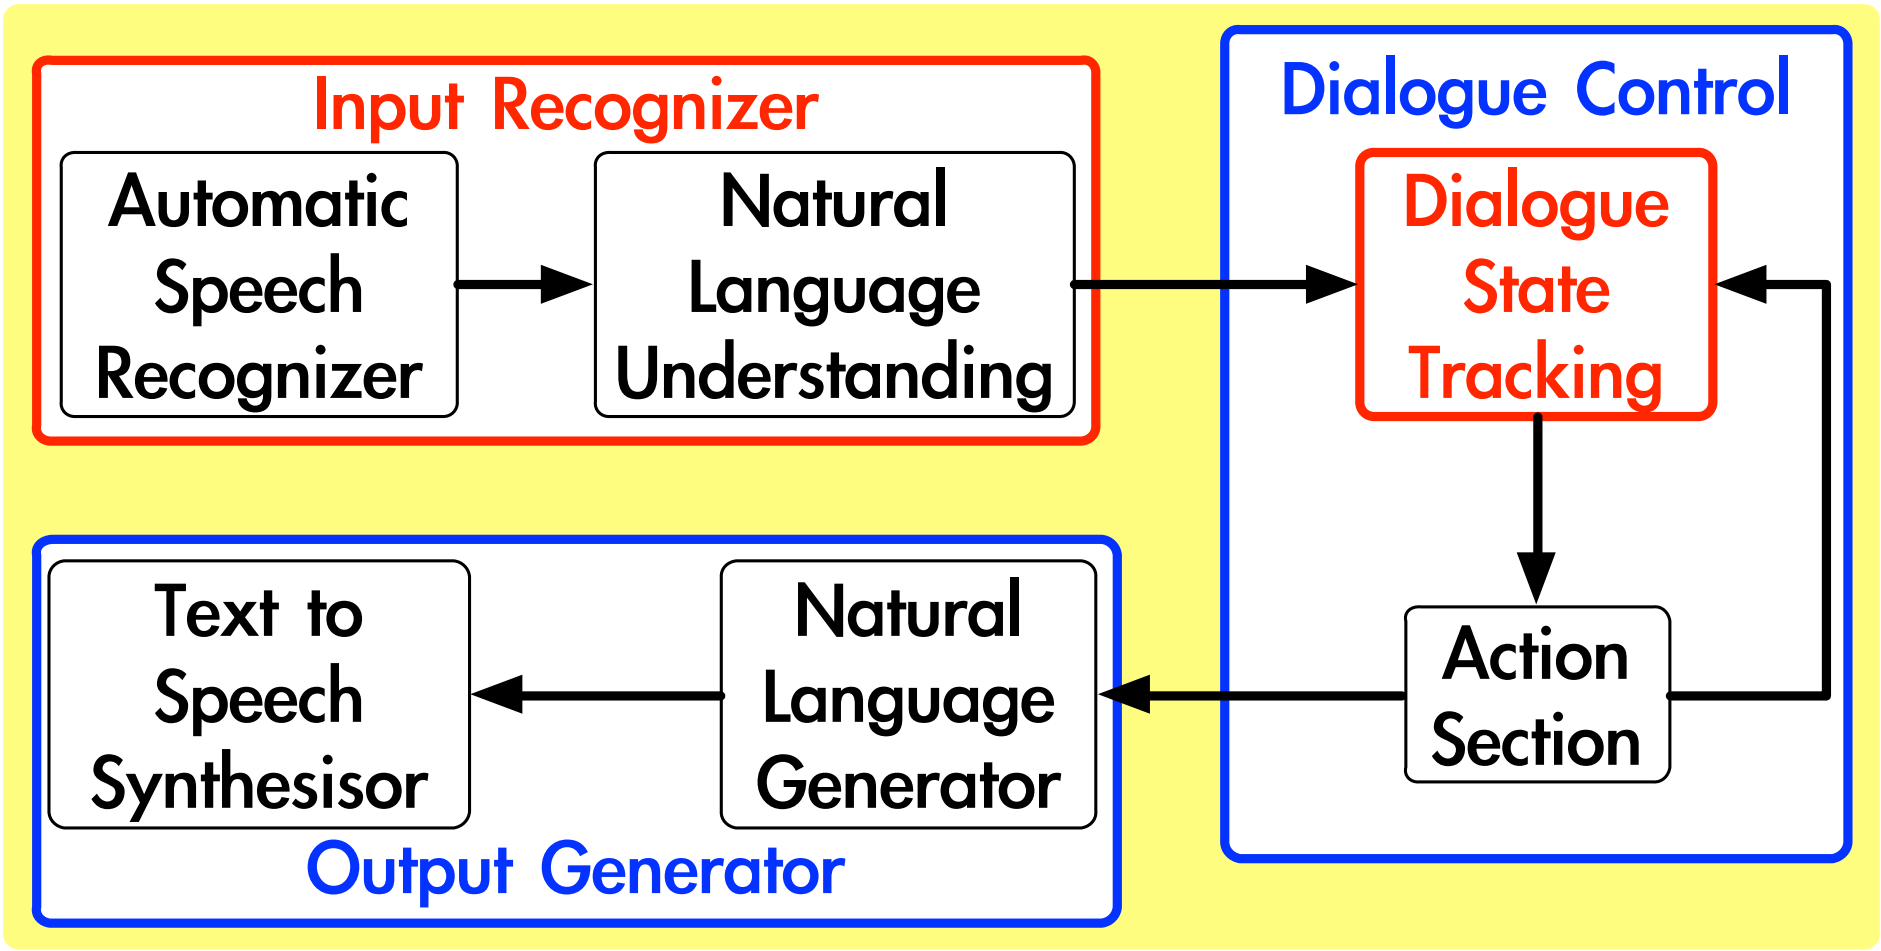

Input Recognition System

feedback with hypothesis to improve recognition accuracy

Dialogue Systems and State Tracking

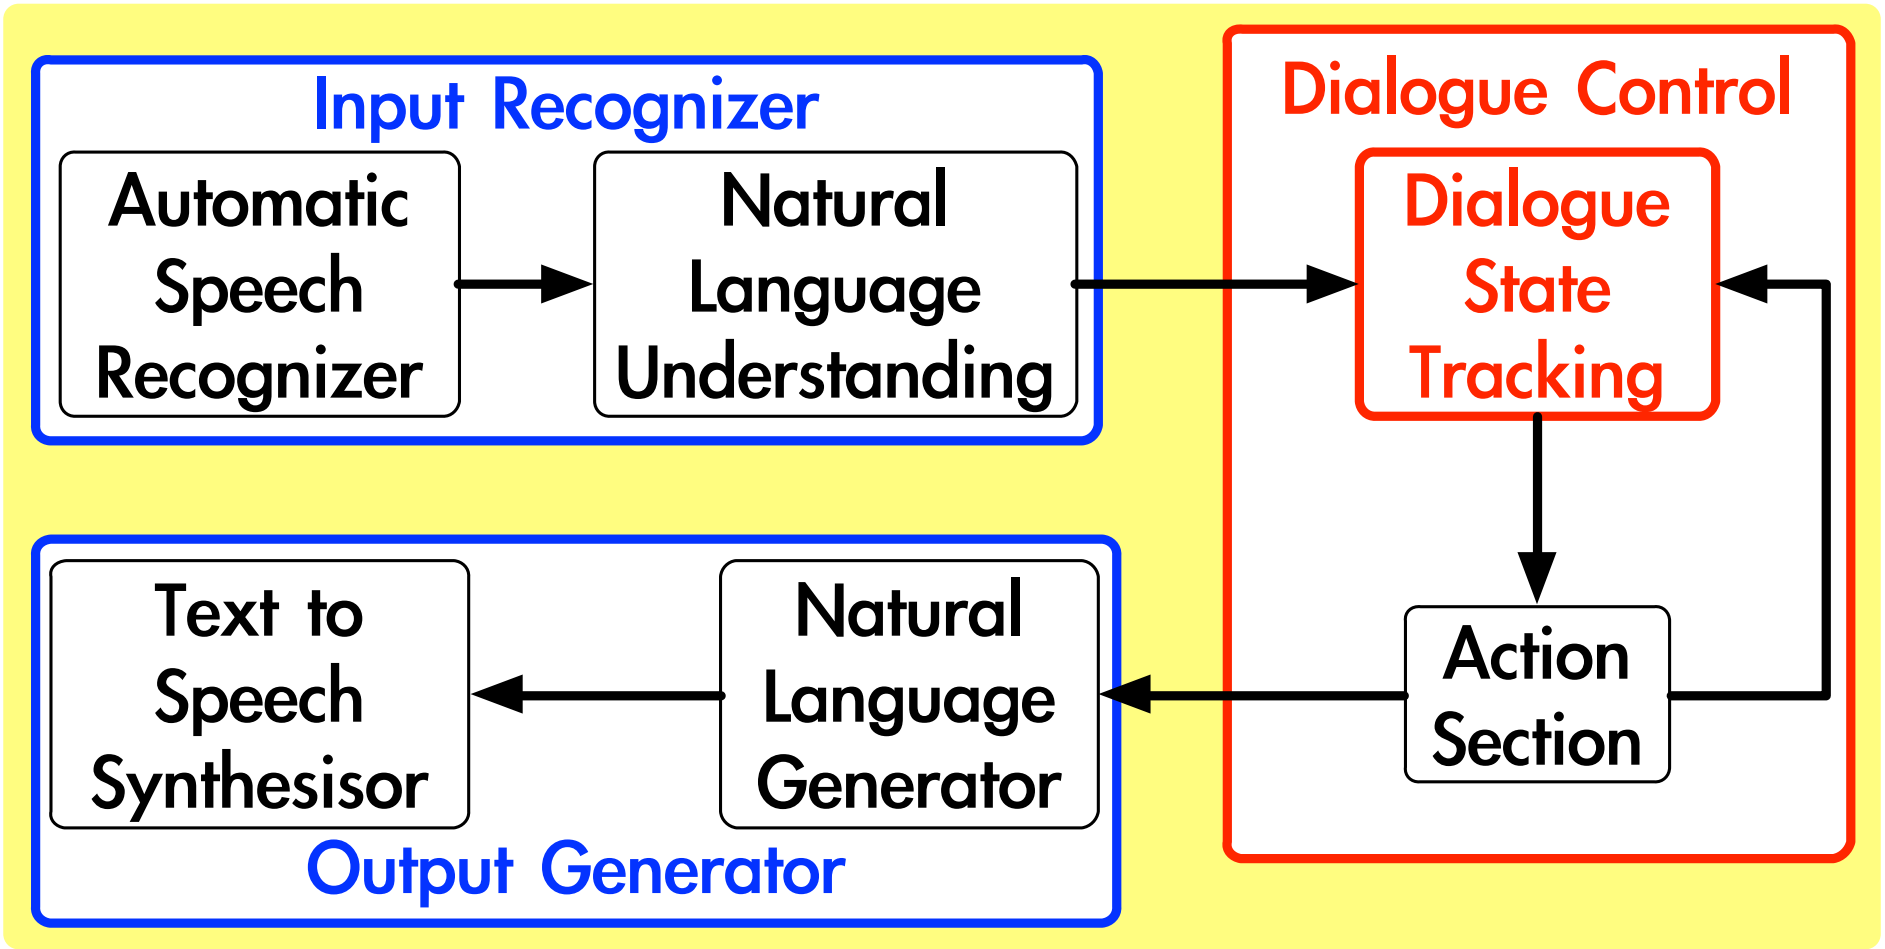

Dialogue Control System

control dialogue flow, for example, reinforcement learning

Dialogue Systems and State Tracking

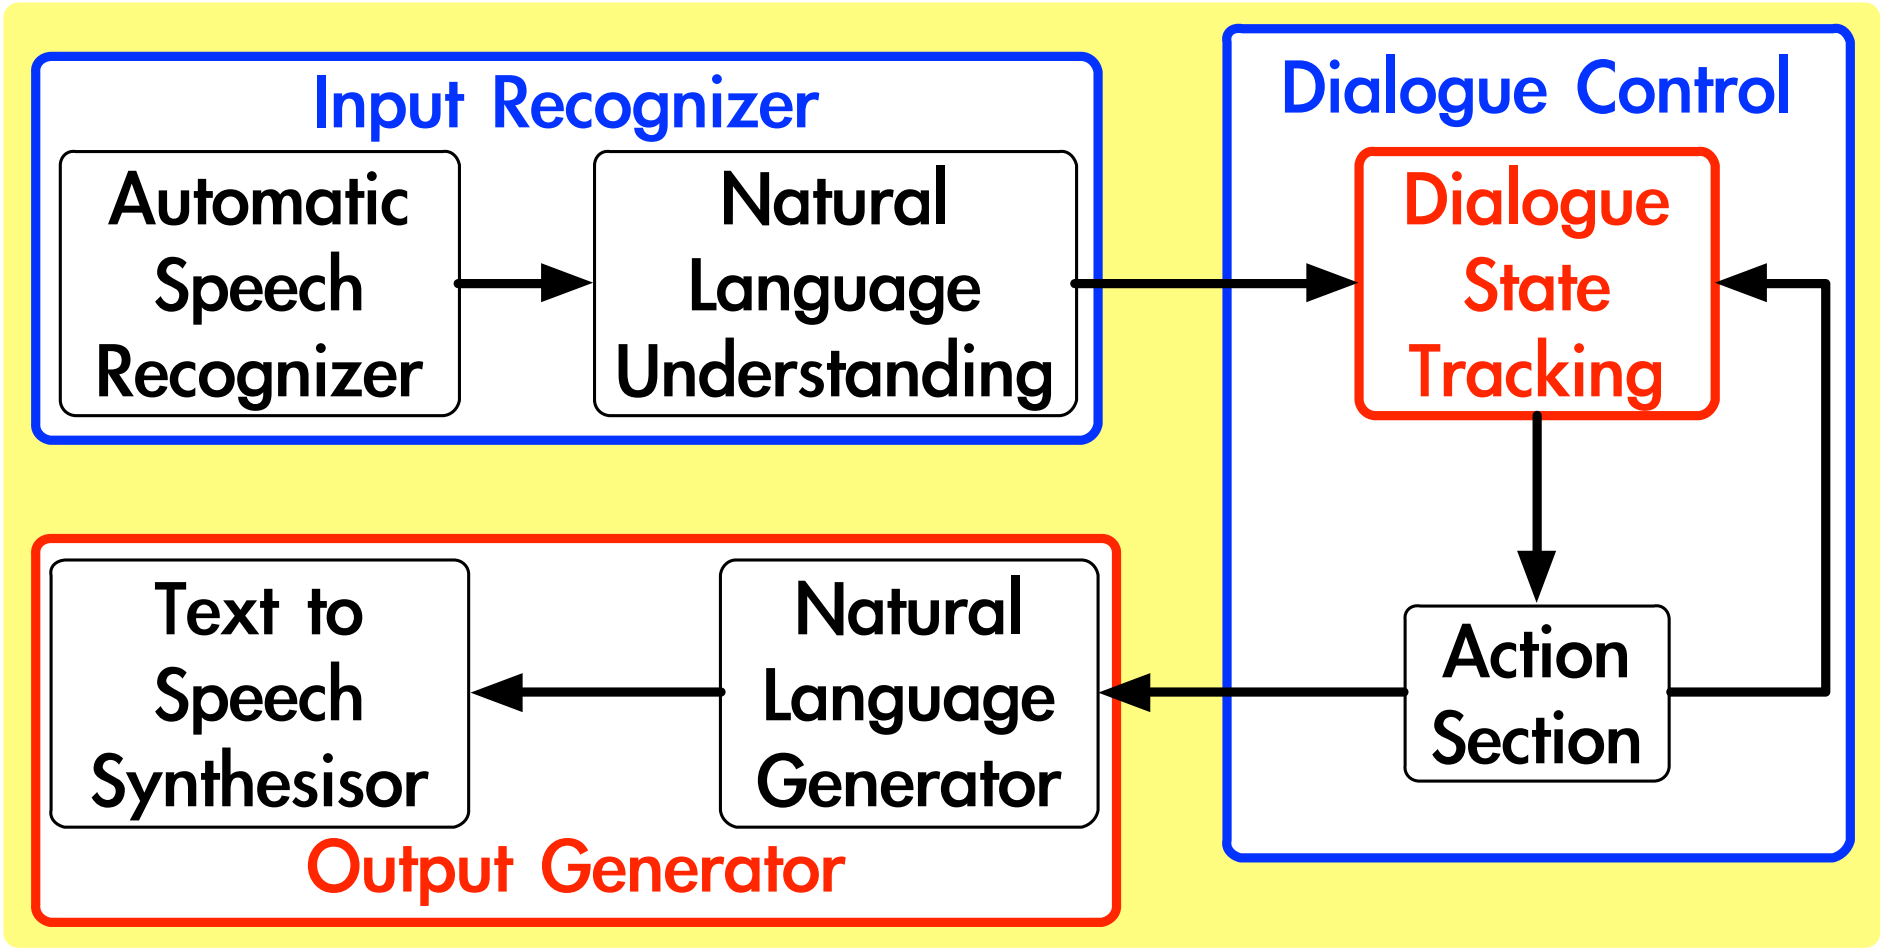

Output Generator System

model languages in conversation and generate meaningful output

Dialogue Systems and State Tracking

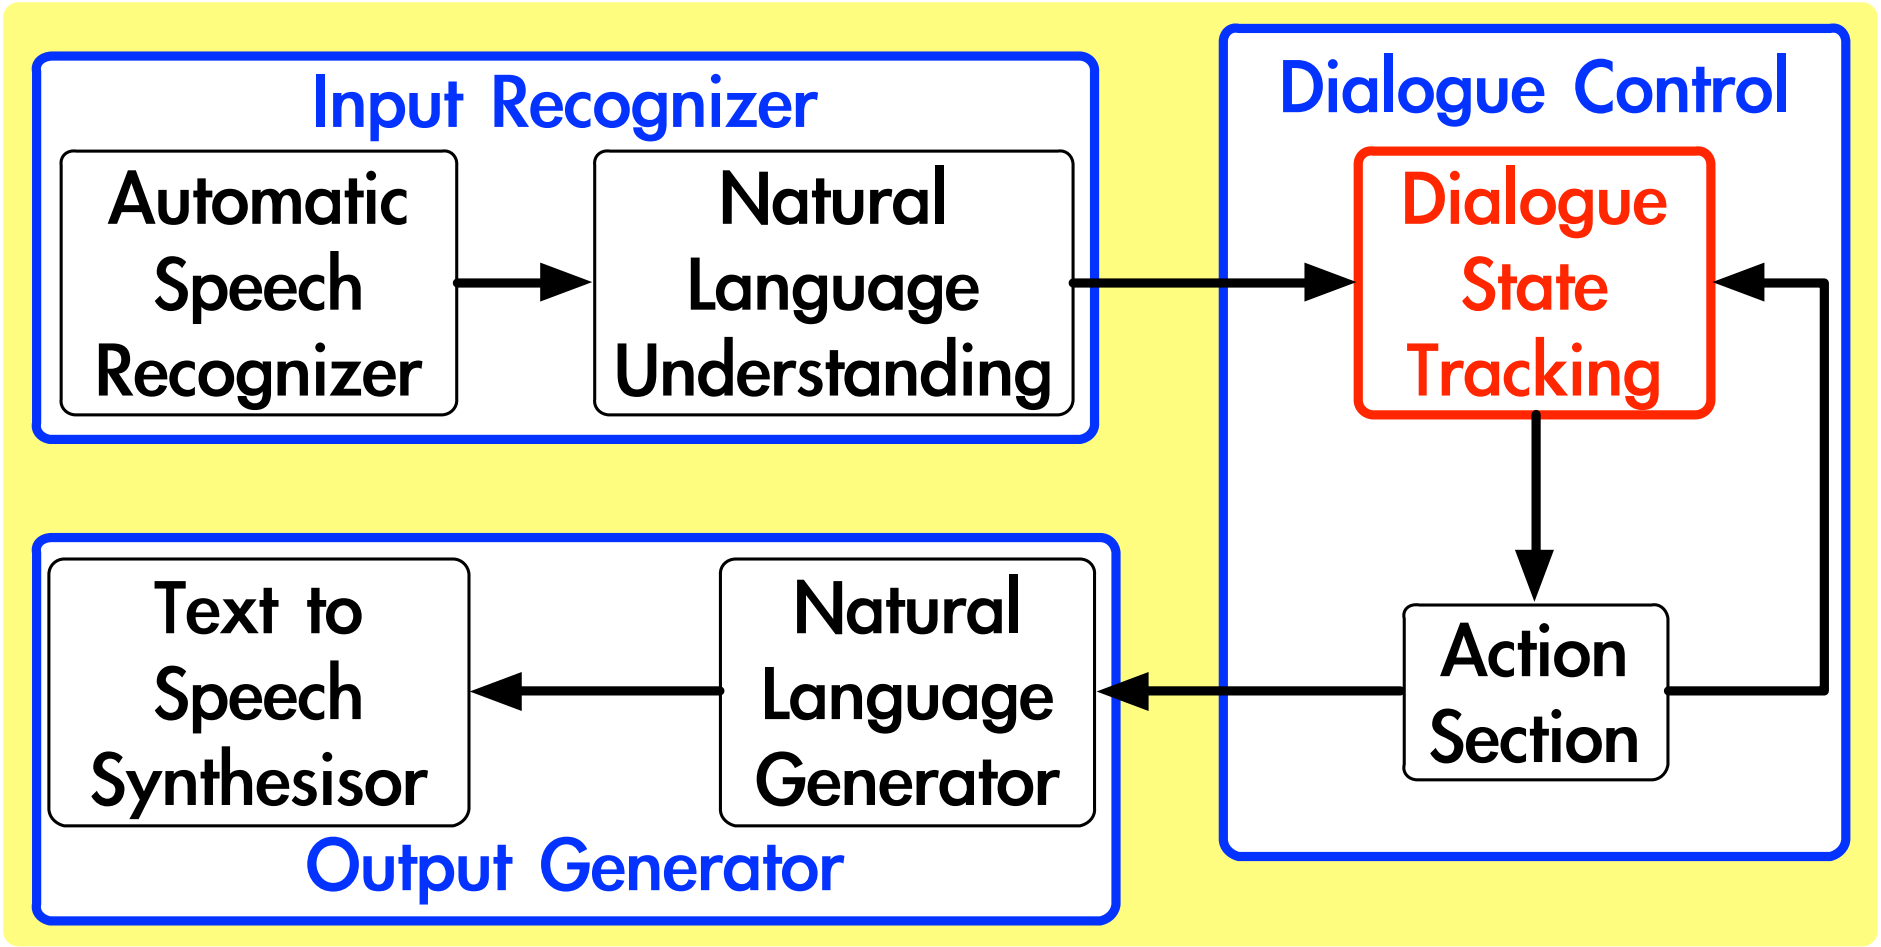

Key Challenge: Dialogue State Tracking

Automatically! Unsupervised! Annotation-free!

Dialogue Systems and State Tracking

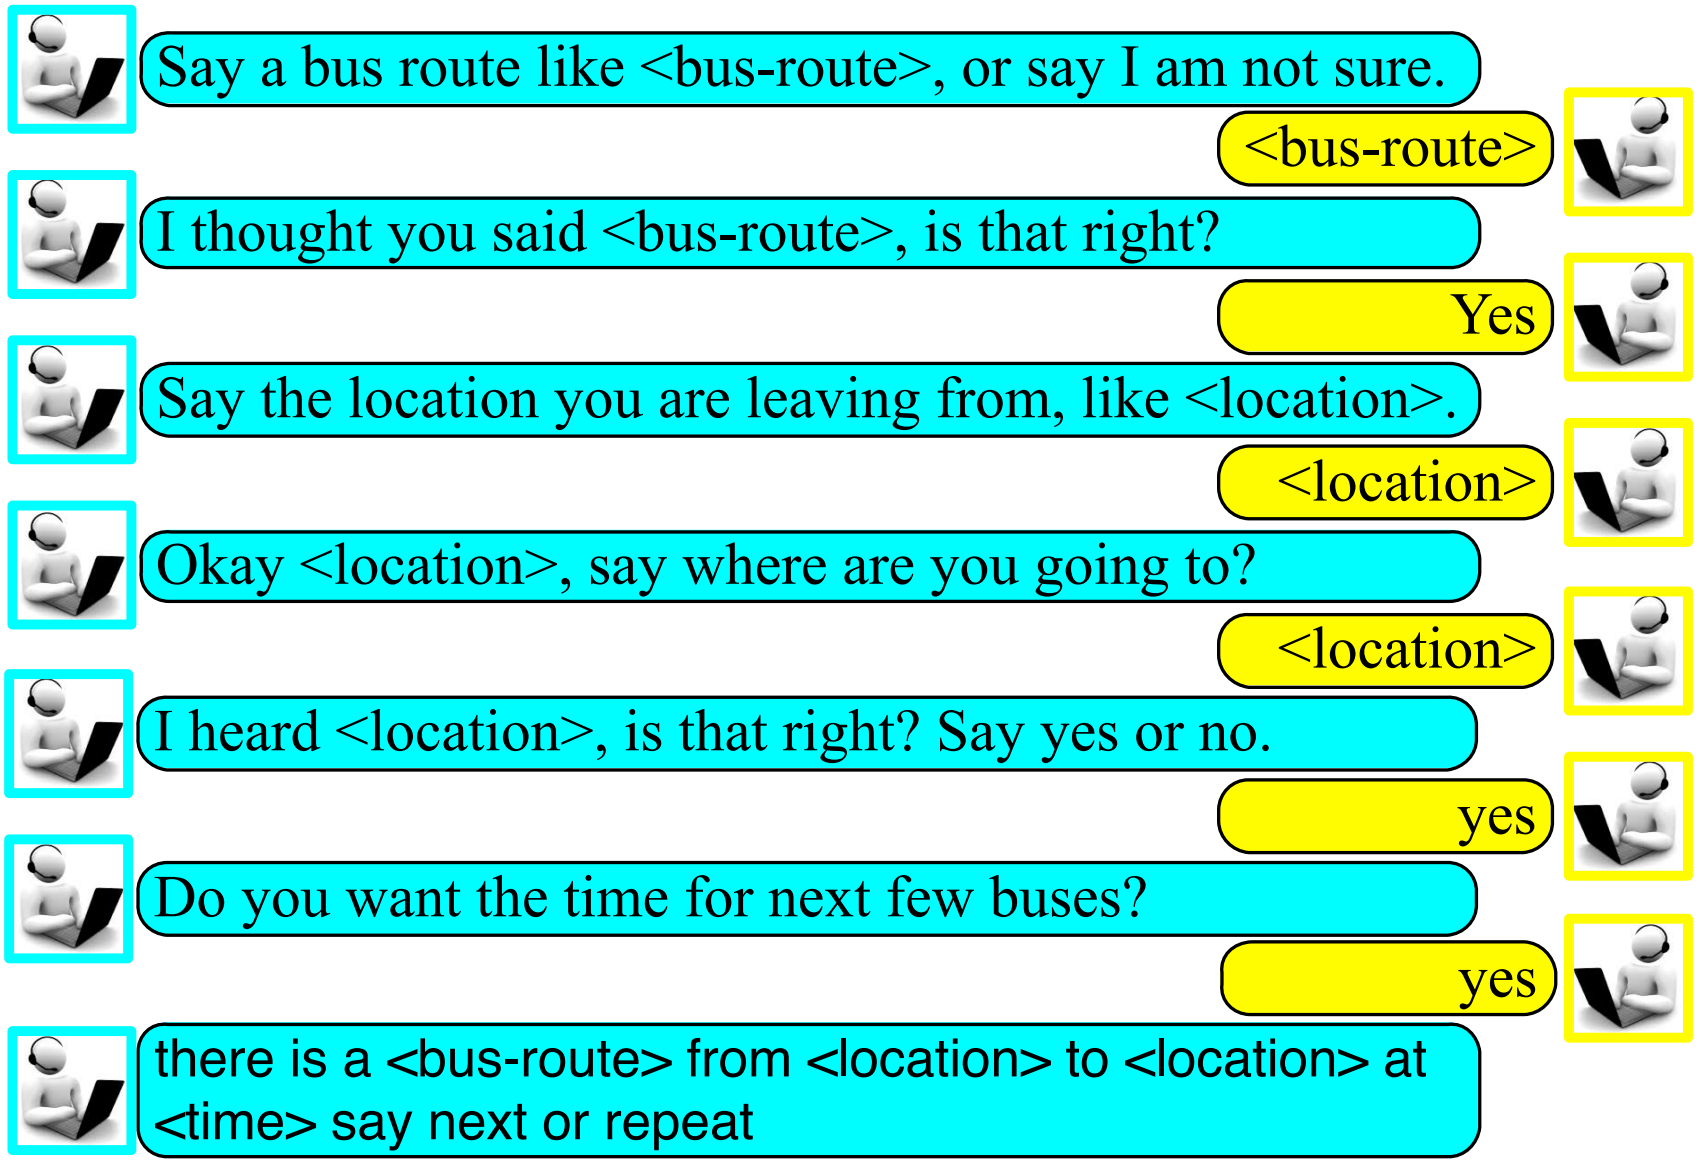

Dialogue Systems and State Tracking

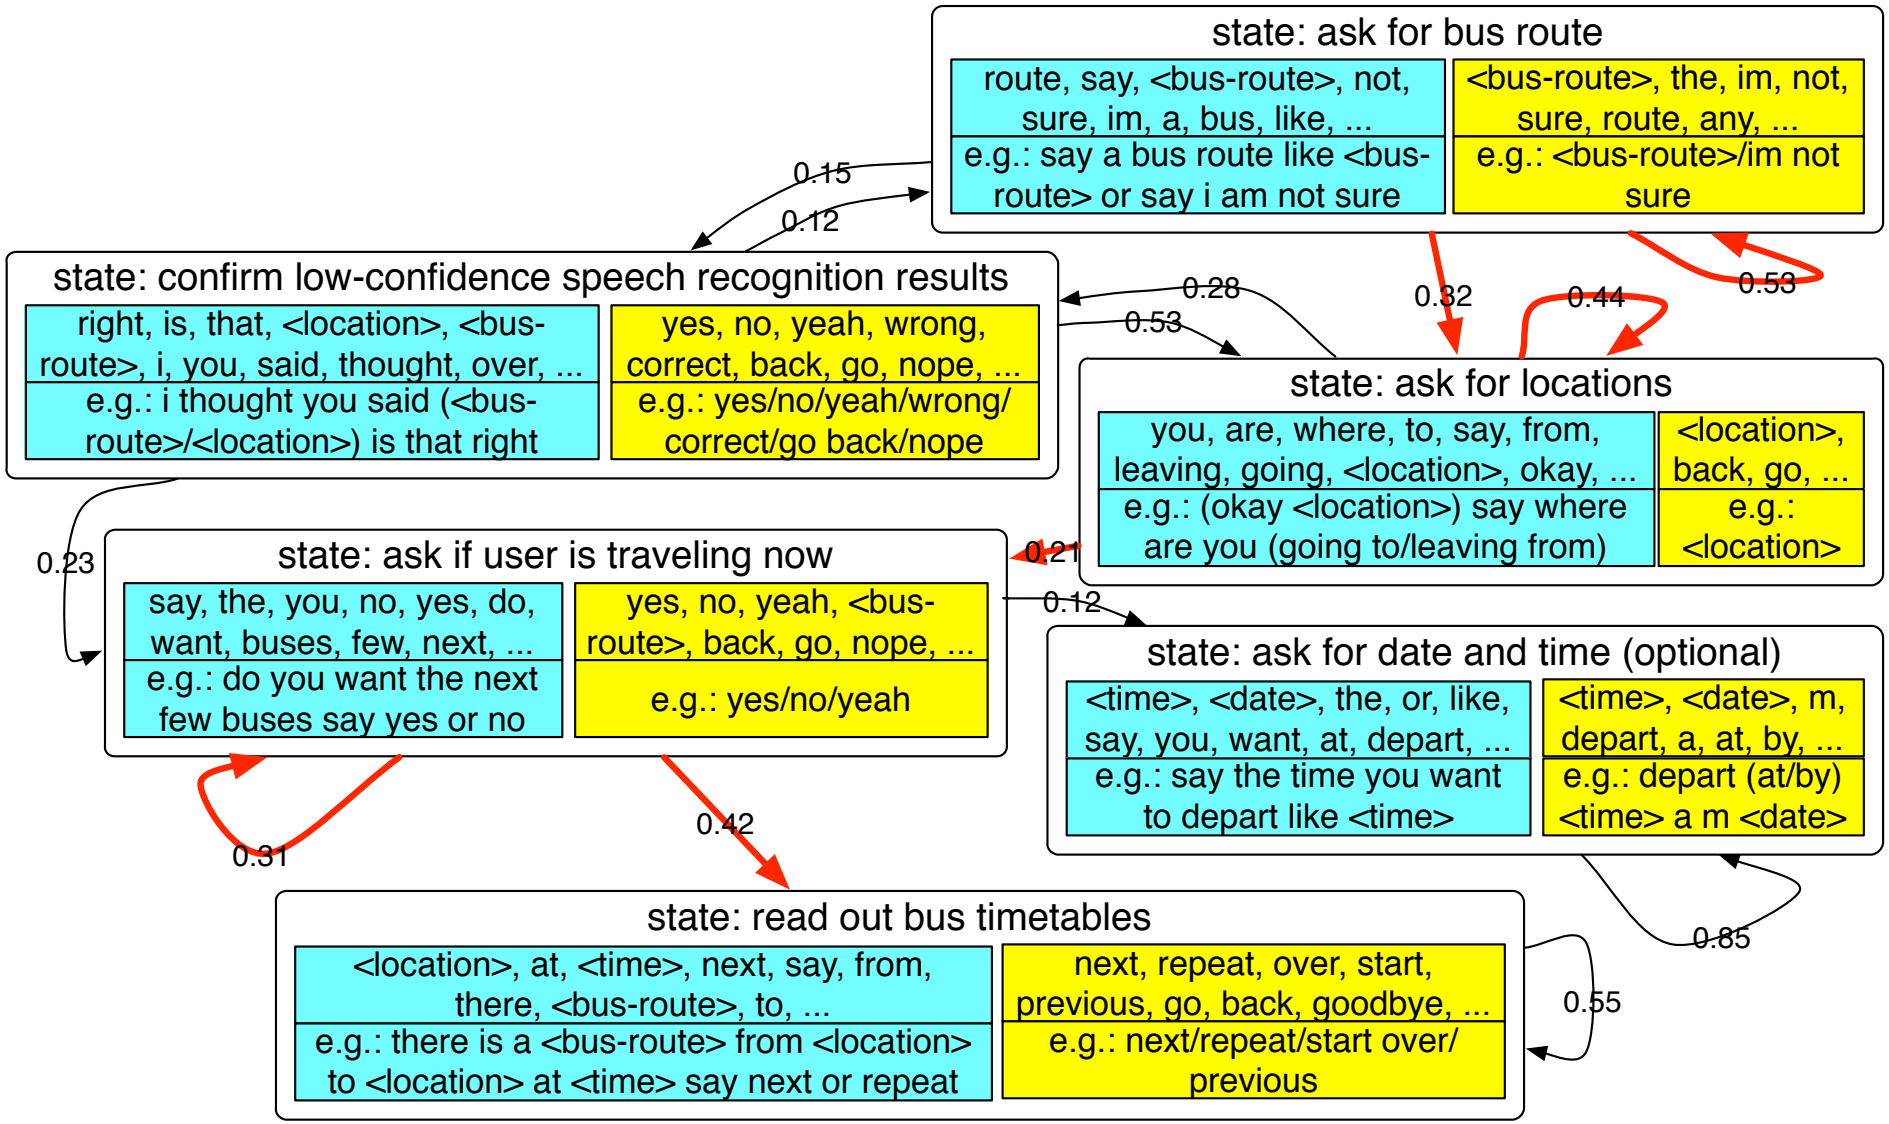

RoadMap

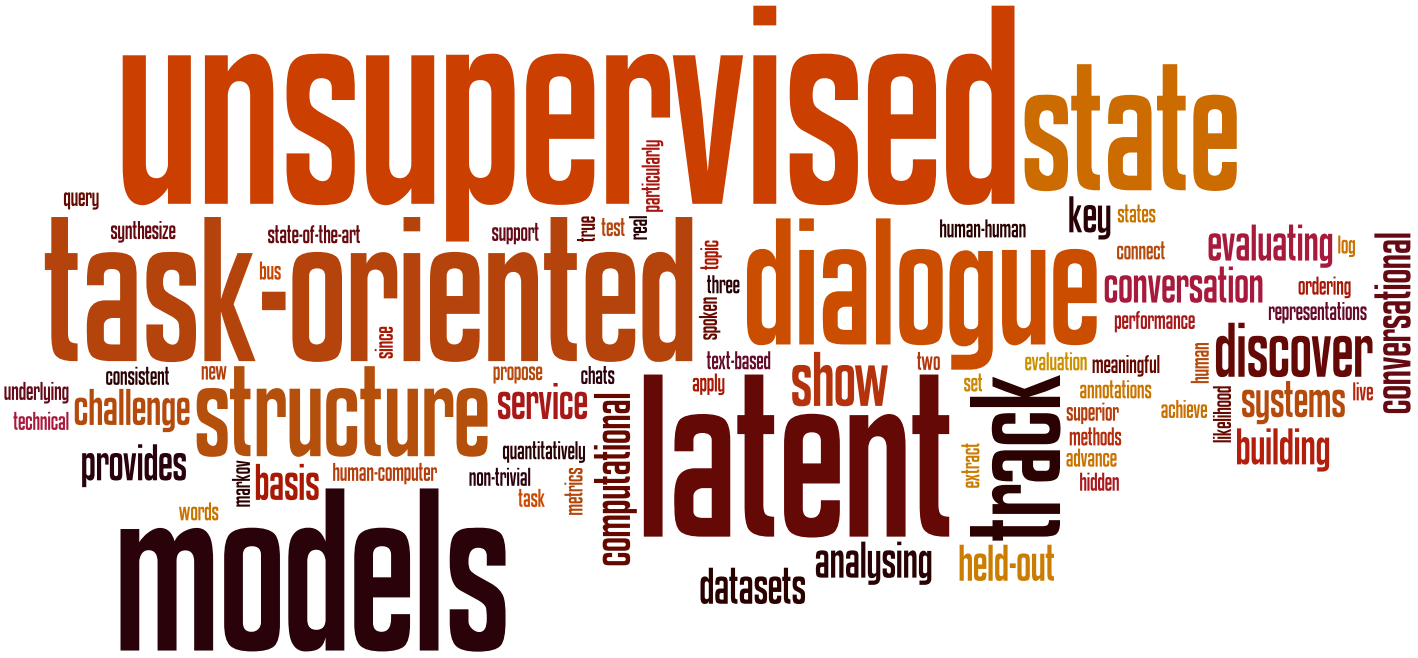

1. Existing Models: LM-HMM and LM-HMMS
2. Proposed Models: TM-HMM and TM-HMMSS
3. Experiments: qualitative and quantitative evaluations
4. Conclusion: what's exciting in Zhai & Williams (2014)...

Outline

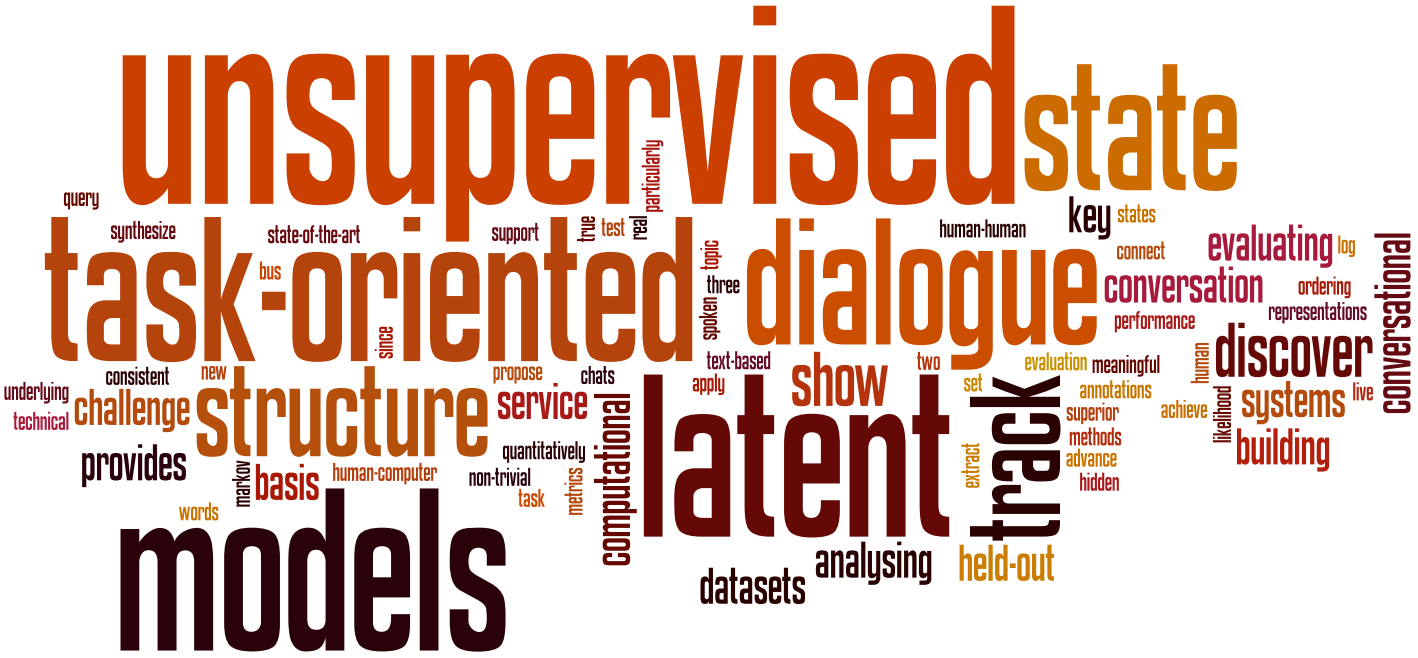

1. Existing Models: LM-HMM and LM-HMMS
2. Proposed Models: TM-HMM and TM-HMMSS
3. Experiments: qualitative and quantitative evaluations
4. Conclusion: what's exciting in Zhai & Williams (2014)...

Motivating Dialogue Example

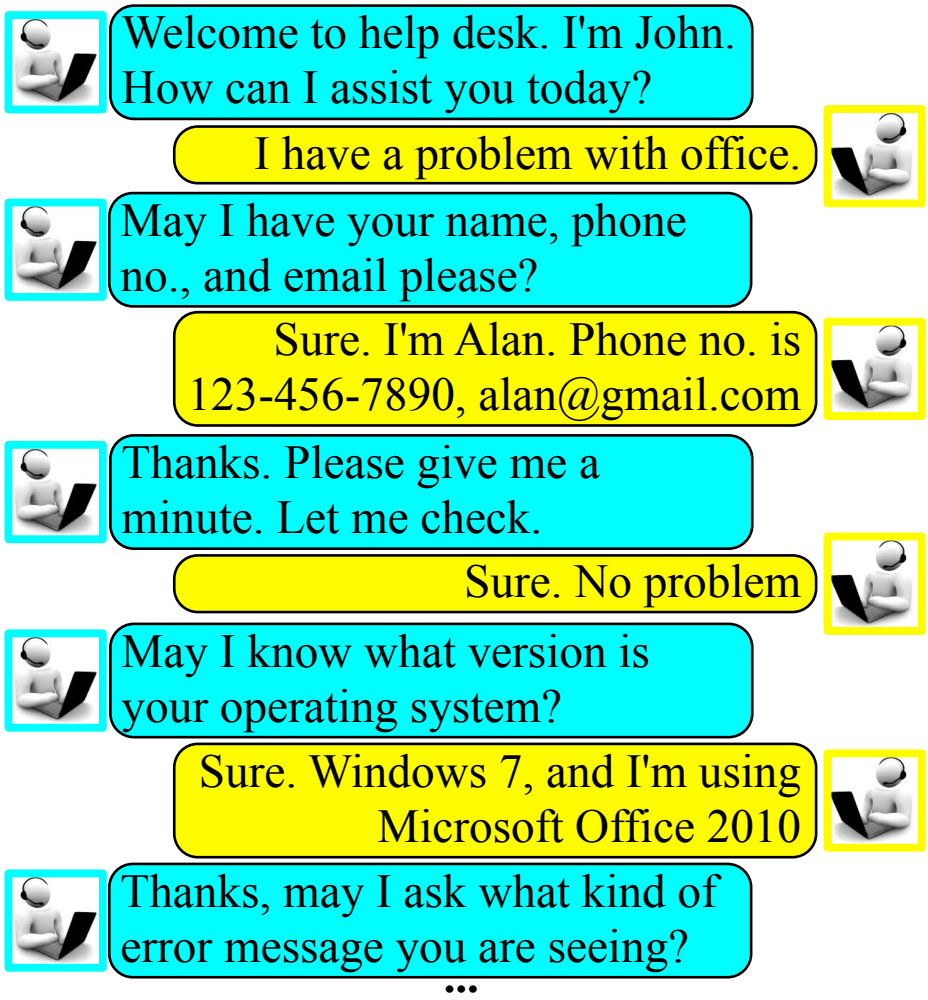

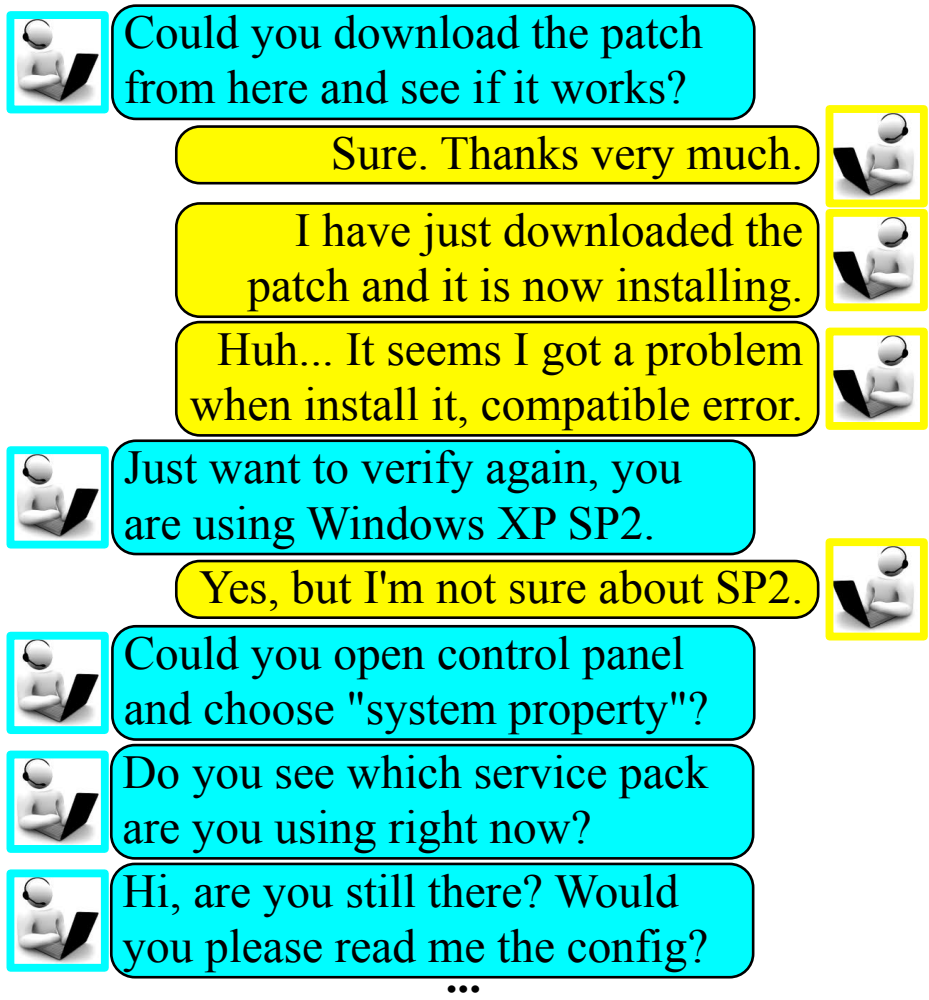

TechSupport: a set of human-human online technical support chats

Conversation Model (Barzilay & Lee, 2004)

LM-HMM: HMM with state language model

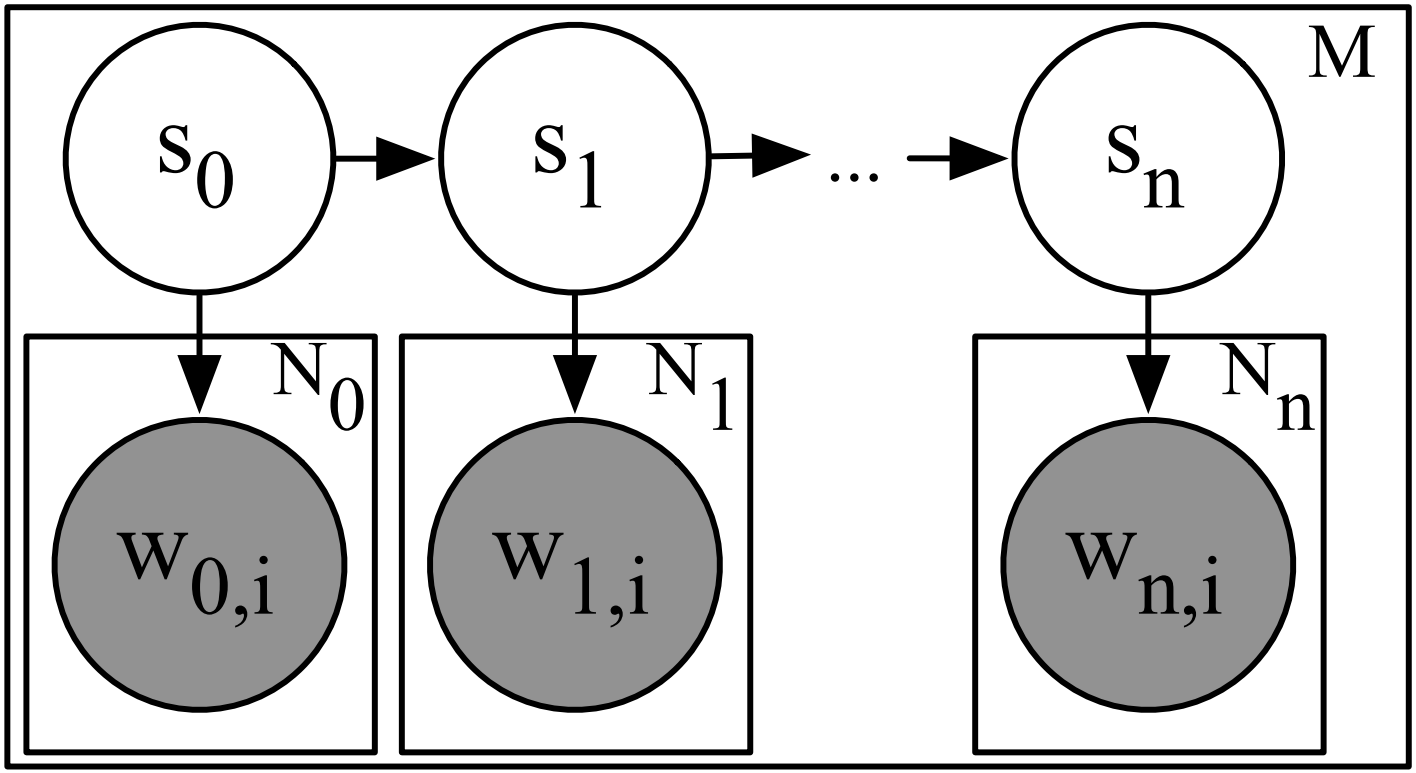

Conversation Model (Barzilay & Lee, 2004)

LM-HMM: HMM with state language model

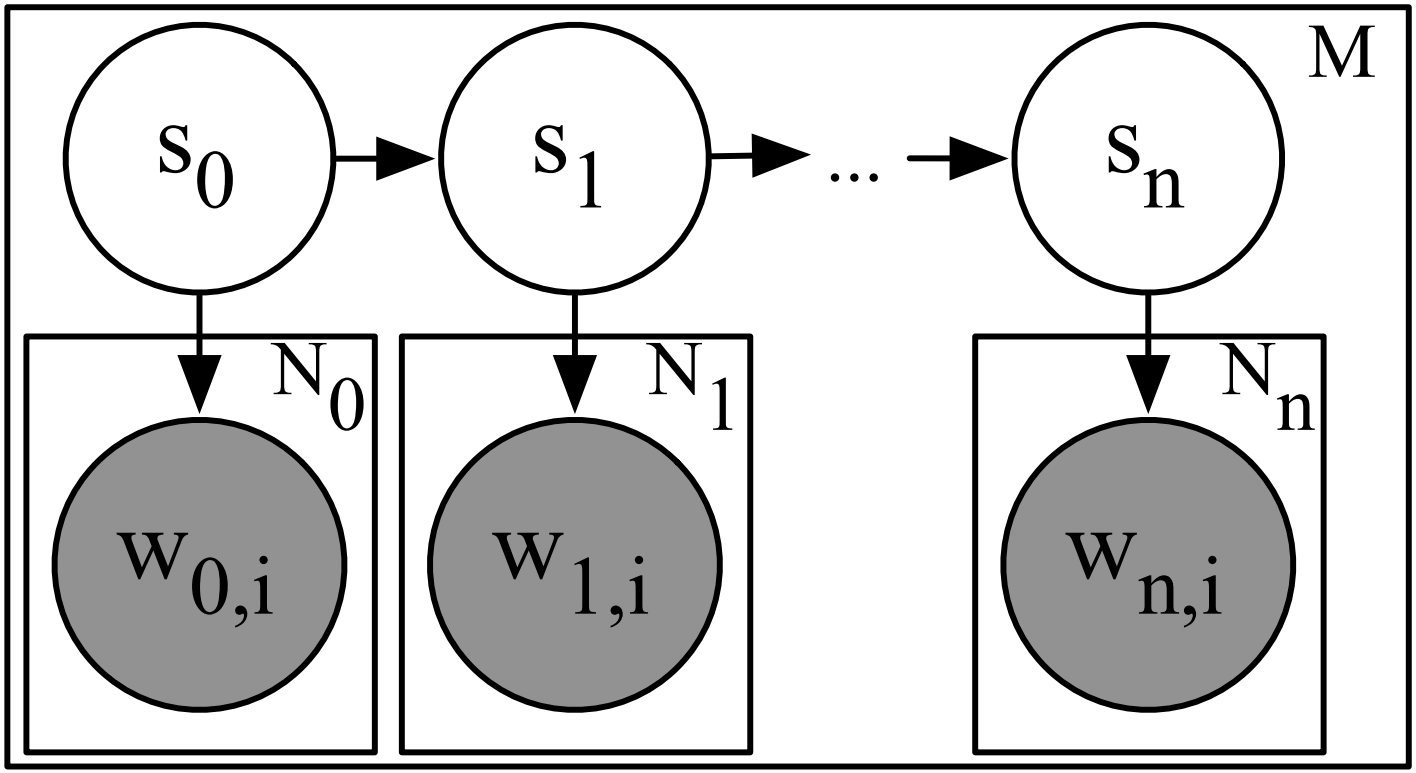

Conversation Model (Barzilay & Lee, 2004)

LM-HMM: HMM with state language model

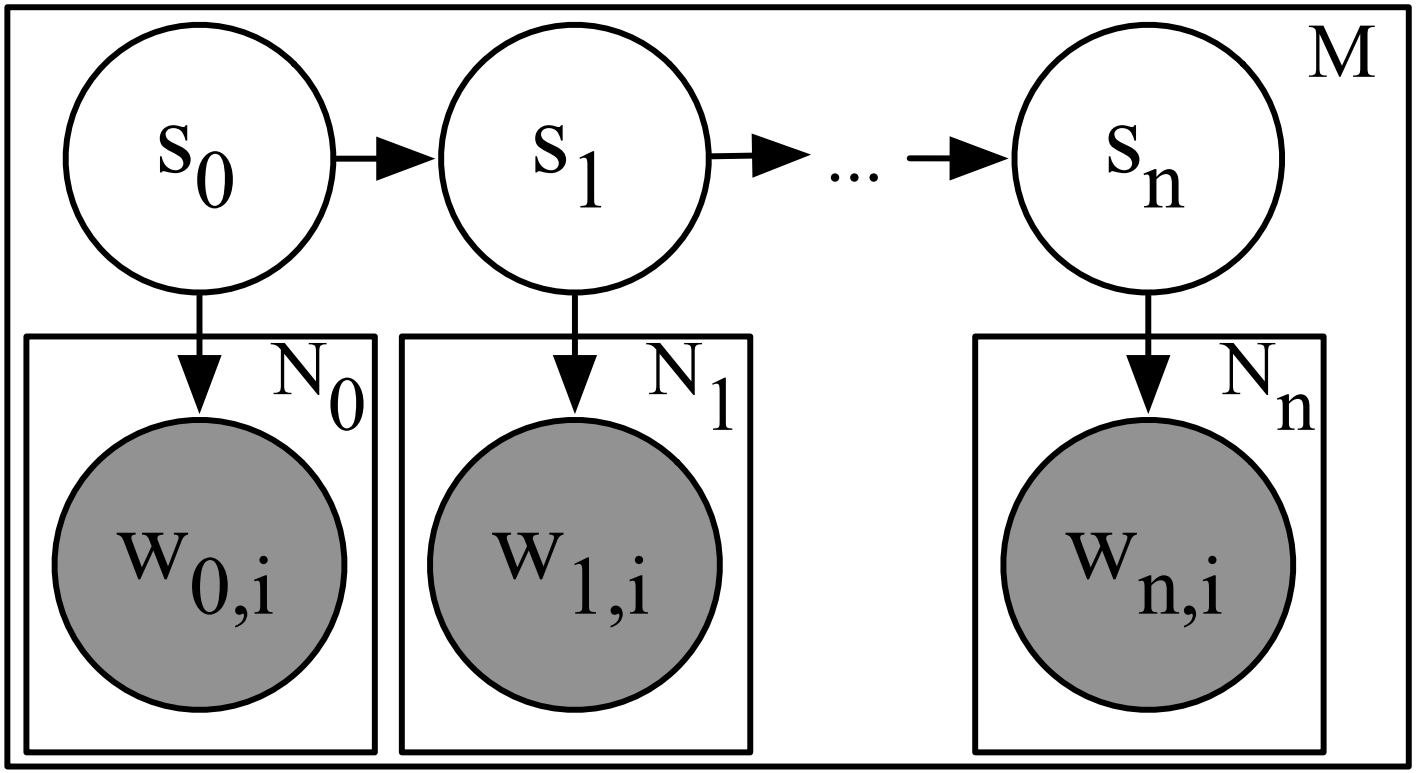

Problem

1. not ALL words are state-dependent.
2. ignores commonalities across different states.

Revisit Motivating Dialogue Example

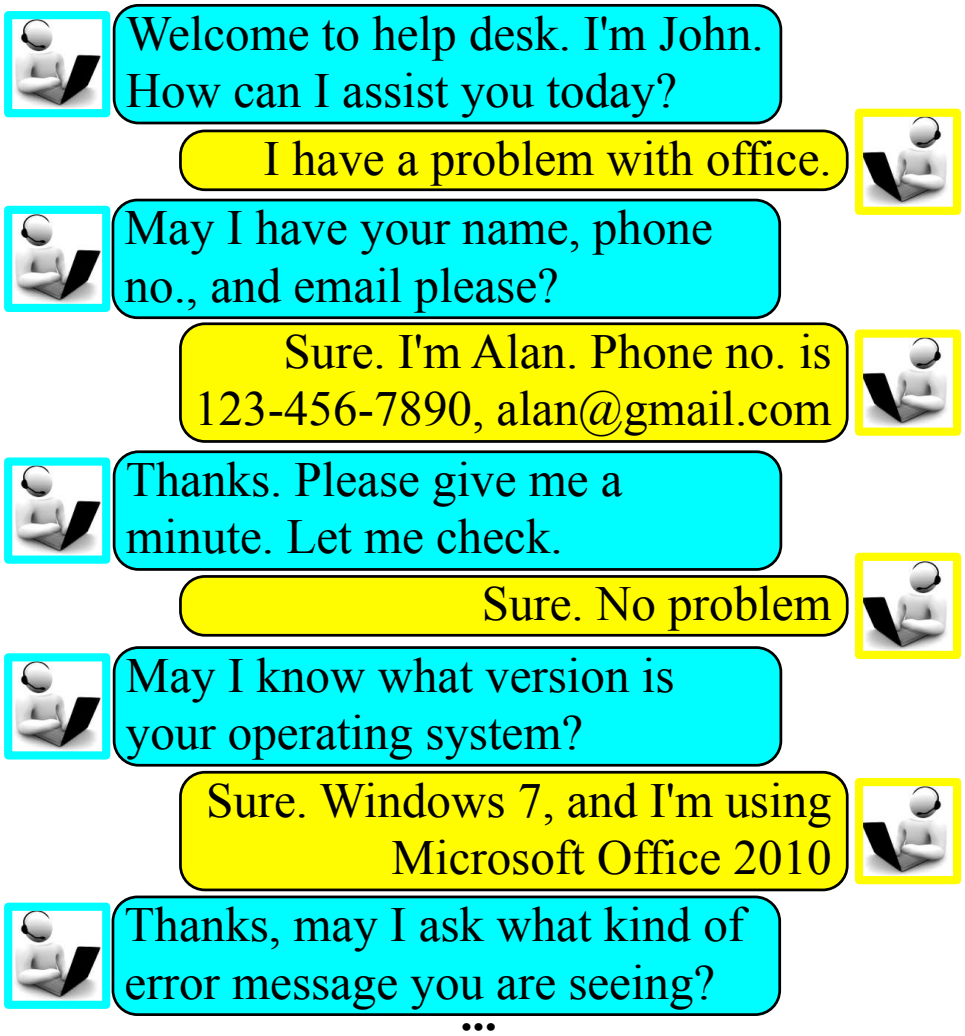

LM-HMM: HMM with state language model

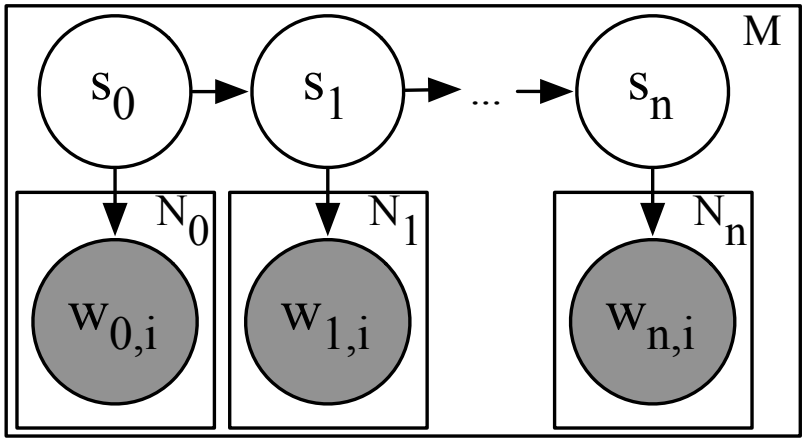

Similar dialogue flow patterns, but *different* problems/domains/topics.

Revisit Motivating Dialogue Example

Welcome to help desk. I'm John. How can I assist you today?

I have a problem with **office**.

May I have your name, phone no., and email please?

Sure. I'm Alan. Phone no. is 123-456-7890, alan@gmail.com

Thanks. Please give me a minute. Let me check.

Sure. No problem

May I know what version is your operating system?

Sure. **Windows 7**, and I'm using **Microsoft Office 2010**

Thanks, may I ask what kind of error message you are seeing?

...

LM-HMM: HMM with state language model

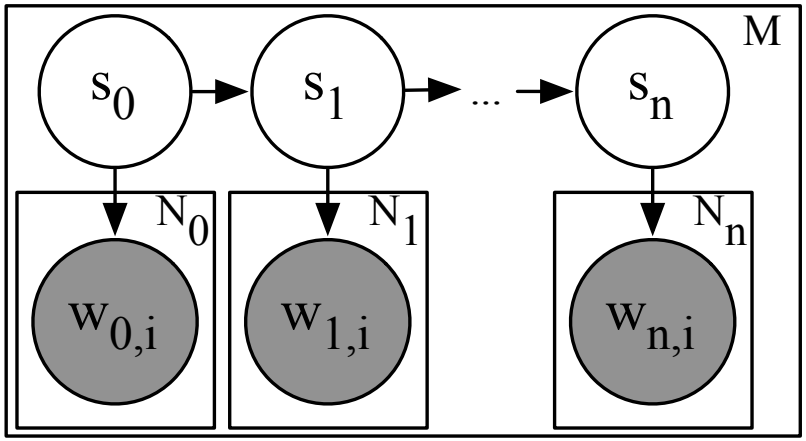

1. not *all* words are state-dependent
2. ignores commonalities across different states/sessions

Similar dialogue flow patterns, but *different* problems/domains/topics.

Conversation + Topic Model (Ritter et al., 2010)

LM-HMMS: HMM with state language model and source generator

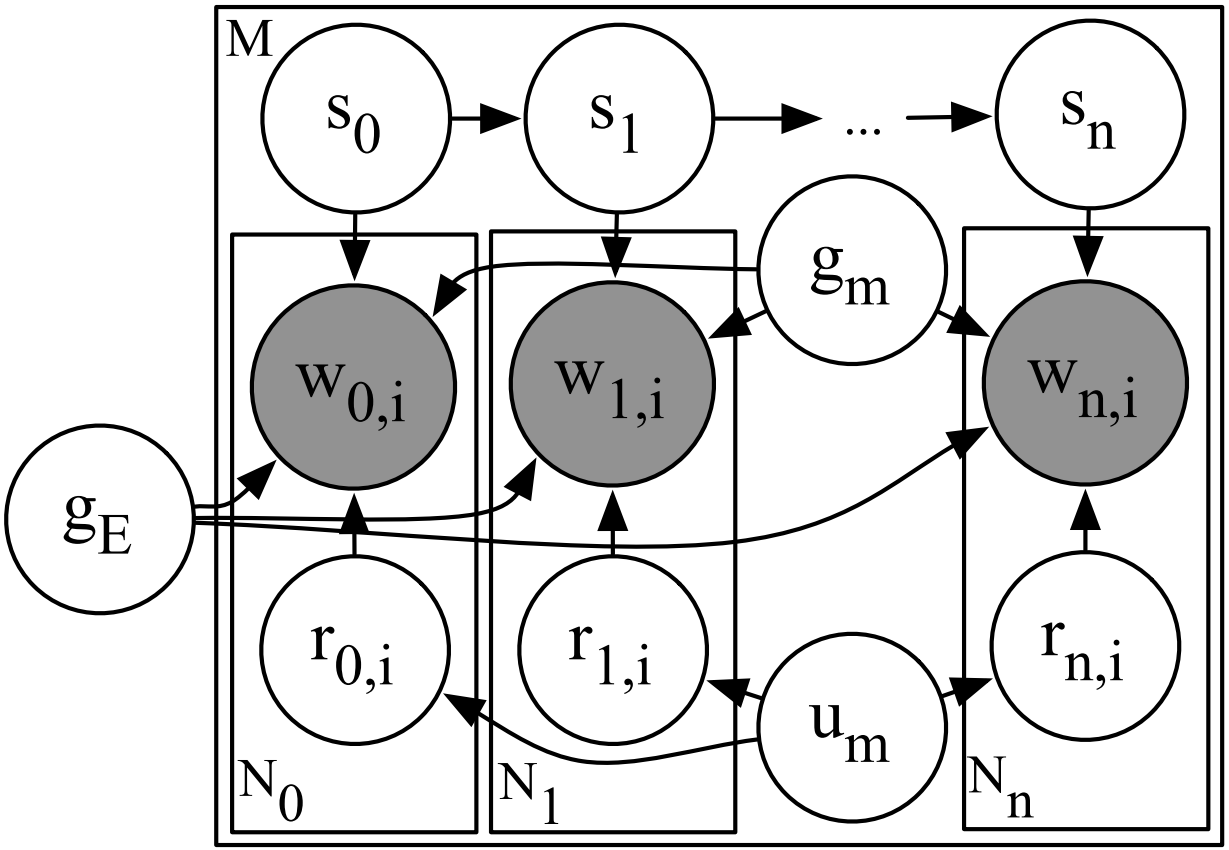

According to source indicator $r_{0,i}$, generate a word from:
1) *state language* model; 2) *session topic* g_m; or 3) background topic g_E.

Conversation + Topic Model (Ritter et al., 2010)

LM-HMMS: HMM with state language model and source generator

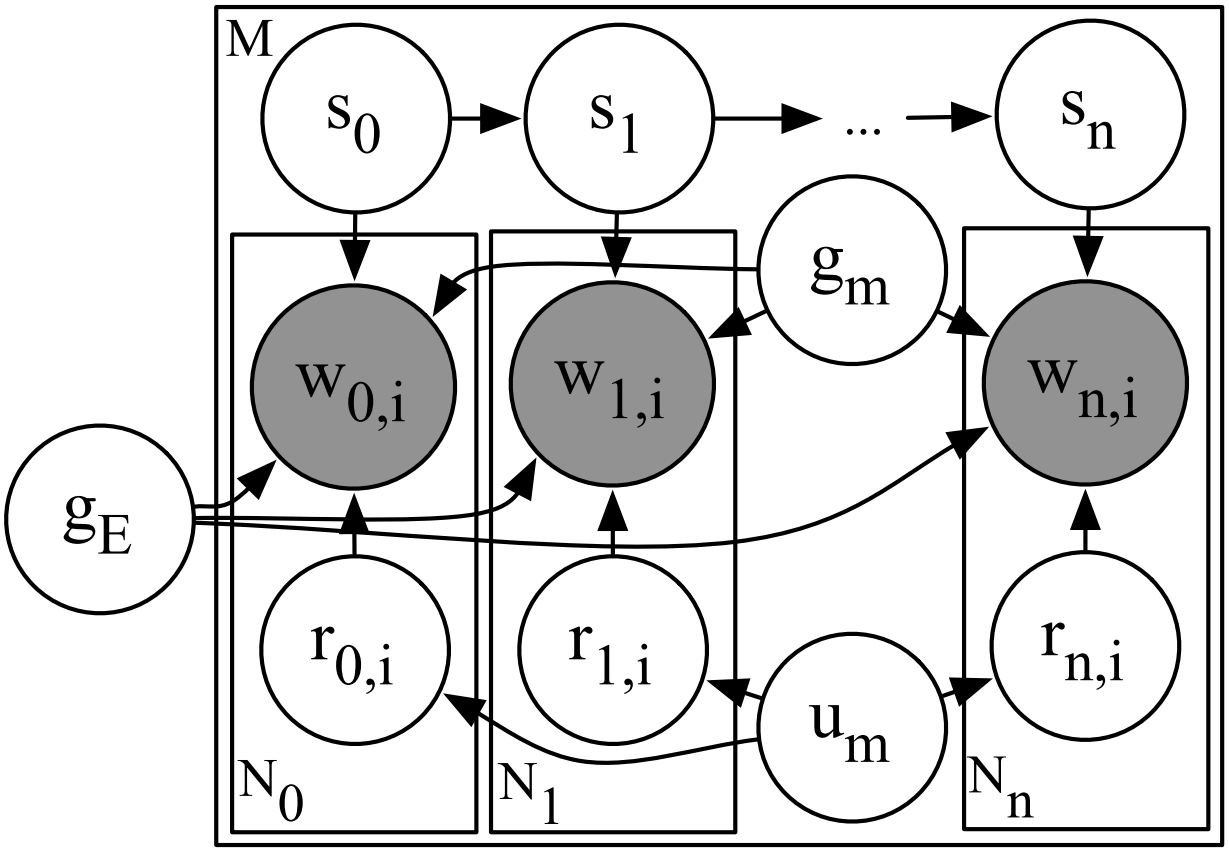

According to source indicator $r_{0,i}$, generate a word from:
1) *state language* model; 2) *session topic* g_m; or 3) background topic g_E.

Conversation + Topic Model (Ritter et al., 2010)

LM-HMMS: HMM with state language model and source generator

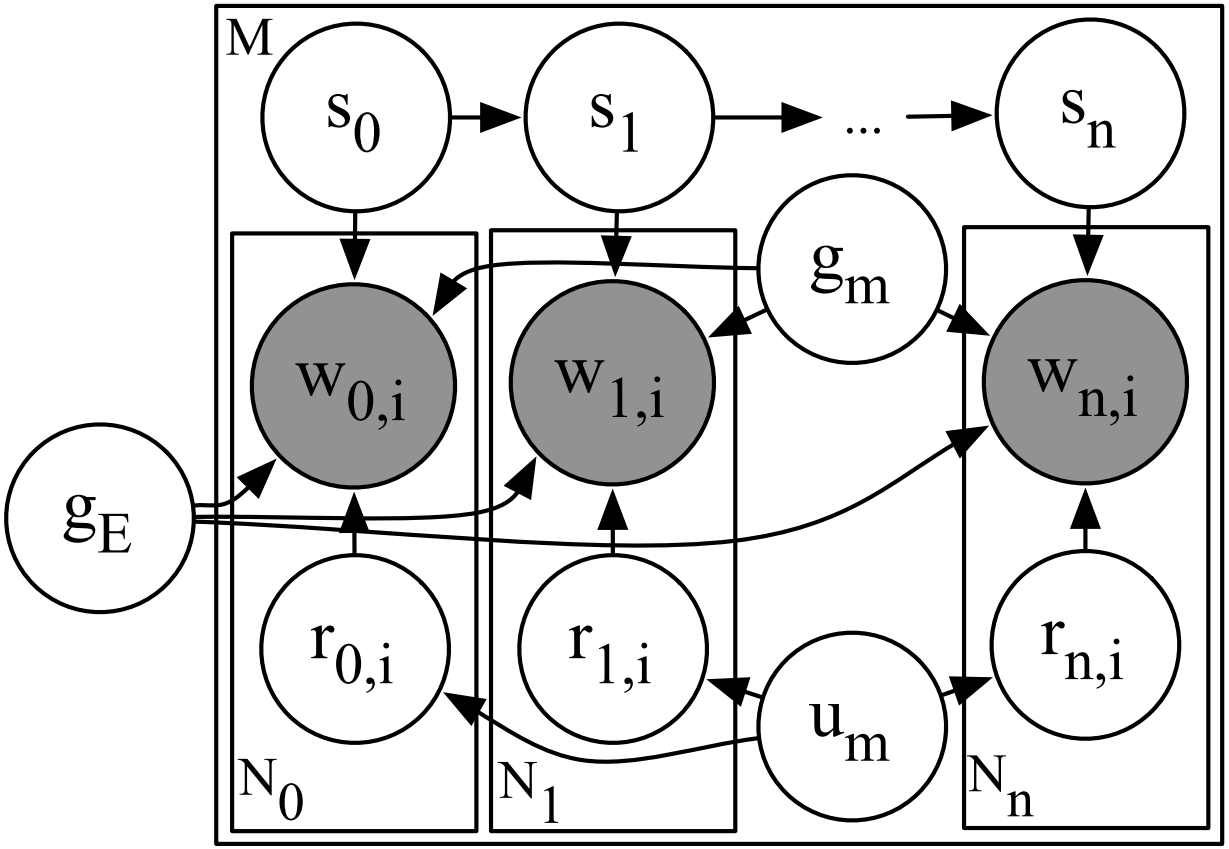

but still, does not share topics across different sessions
fine for social media dialogues, e.g. twitter chats
but not so well for task-oriented dialogues

Free-style Dialogues vs. Task-oriented Dialogues

Free-style
Diffused Topics

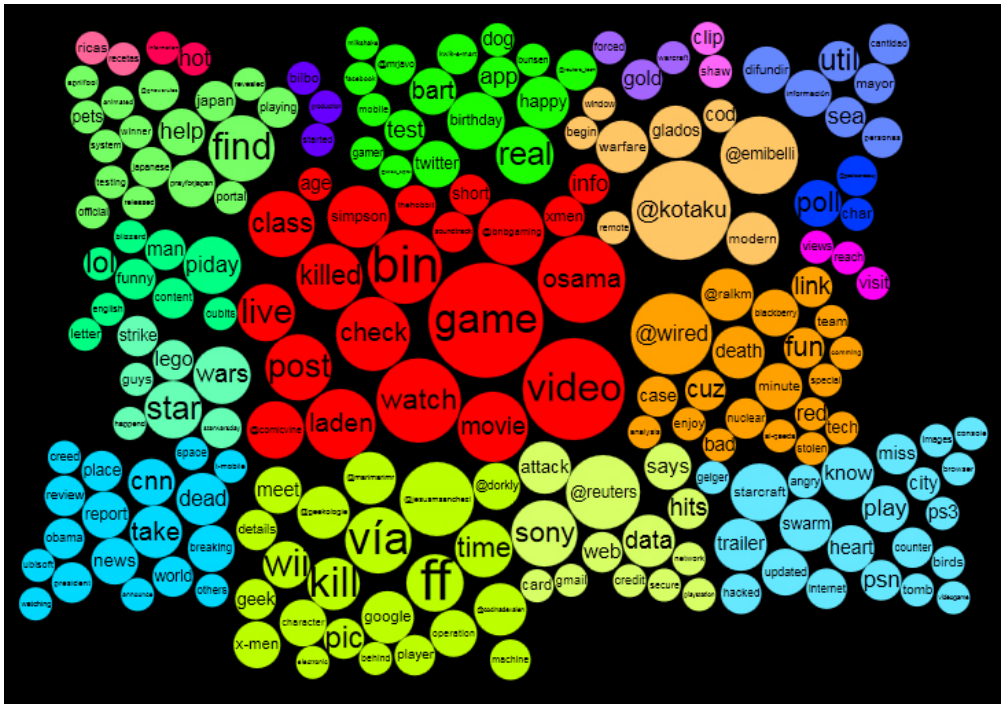

Free-style Dialogues vs. Task-oriented Dialogues

Free-style
Diffused Topics

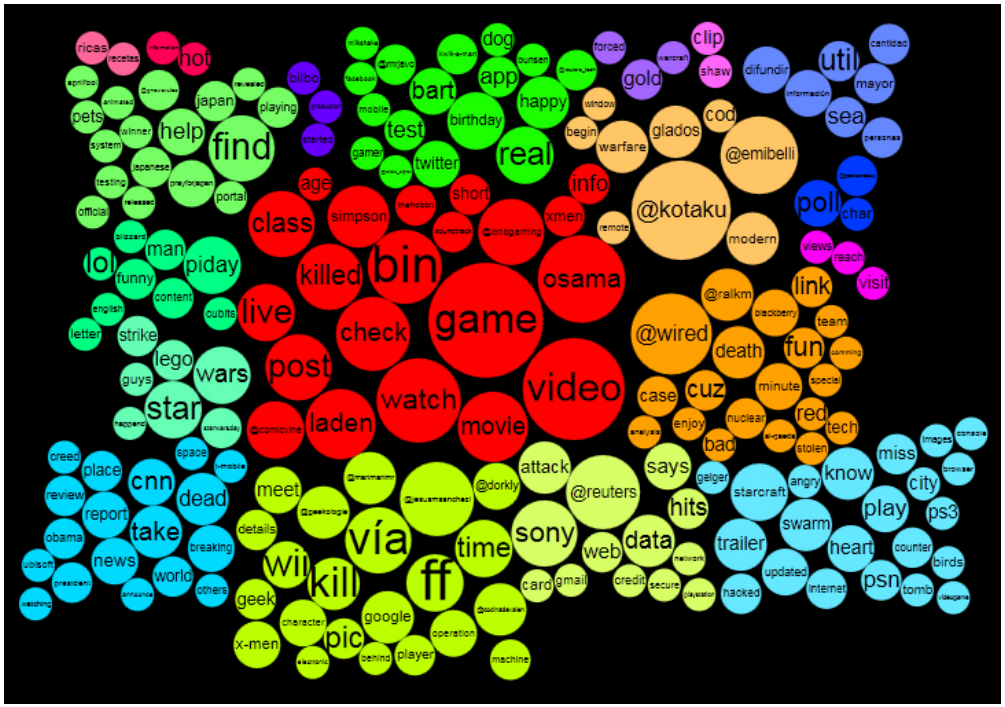

Task-oriented
In-domain Topics

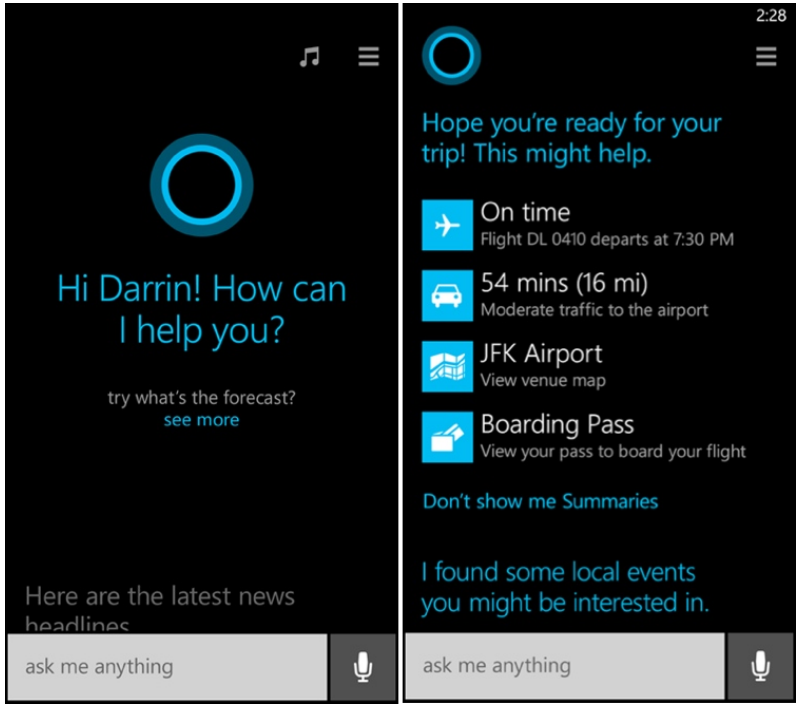

Free-style Dialogues vs. Task-oriented Dialogues

Free-style
Diffused Topics

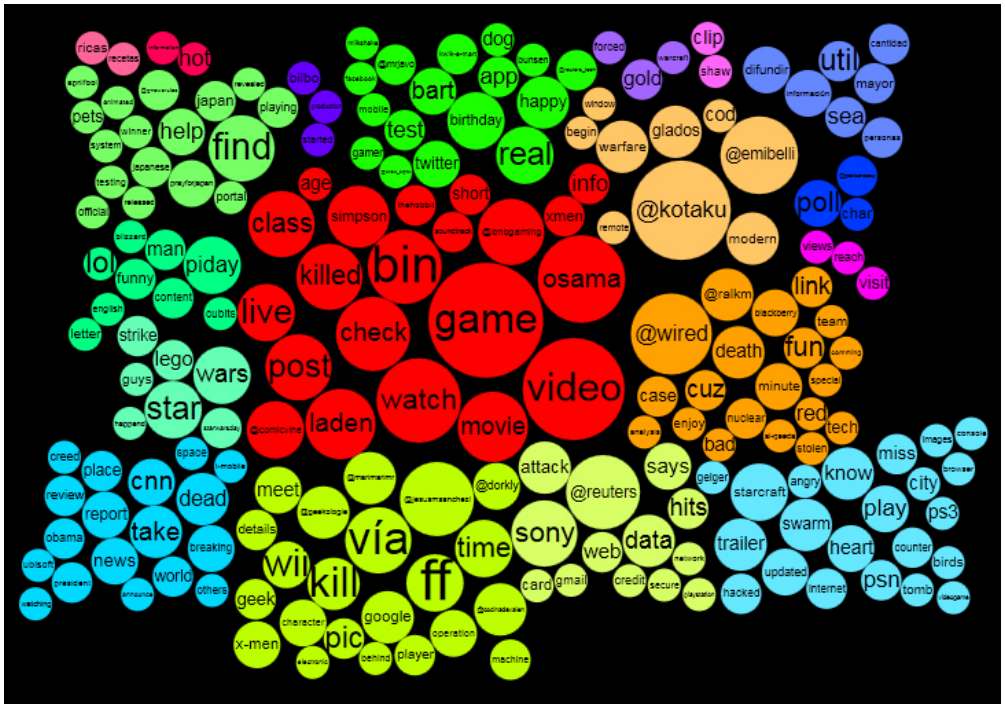

Task-oriented
In-domain Topics

Task-oriented Dialogues

Topics are shared across different dialogues, or maybe even across different utterances in a dialogue session.

Outline

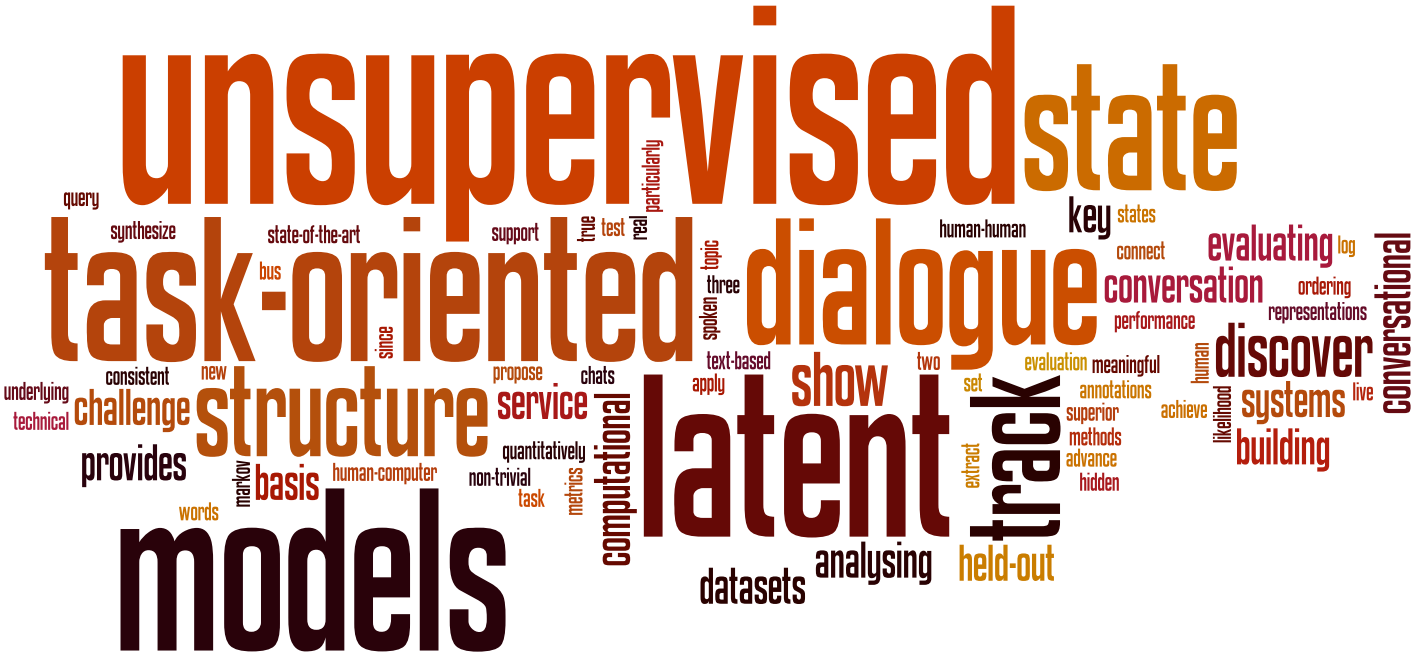

1. Existing Models: LM-HMM and LM-HMMS
2. Proposed Models: TM-HMM and TM-HMMSS
3. Experiments: qualitative and quantitative evaluations
4. Conclusion: what's exciting in Zhai & Williams (2014)...

TM-HMM: Share Topics across States

LM-HMM: HMM with state language model

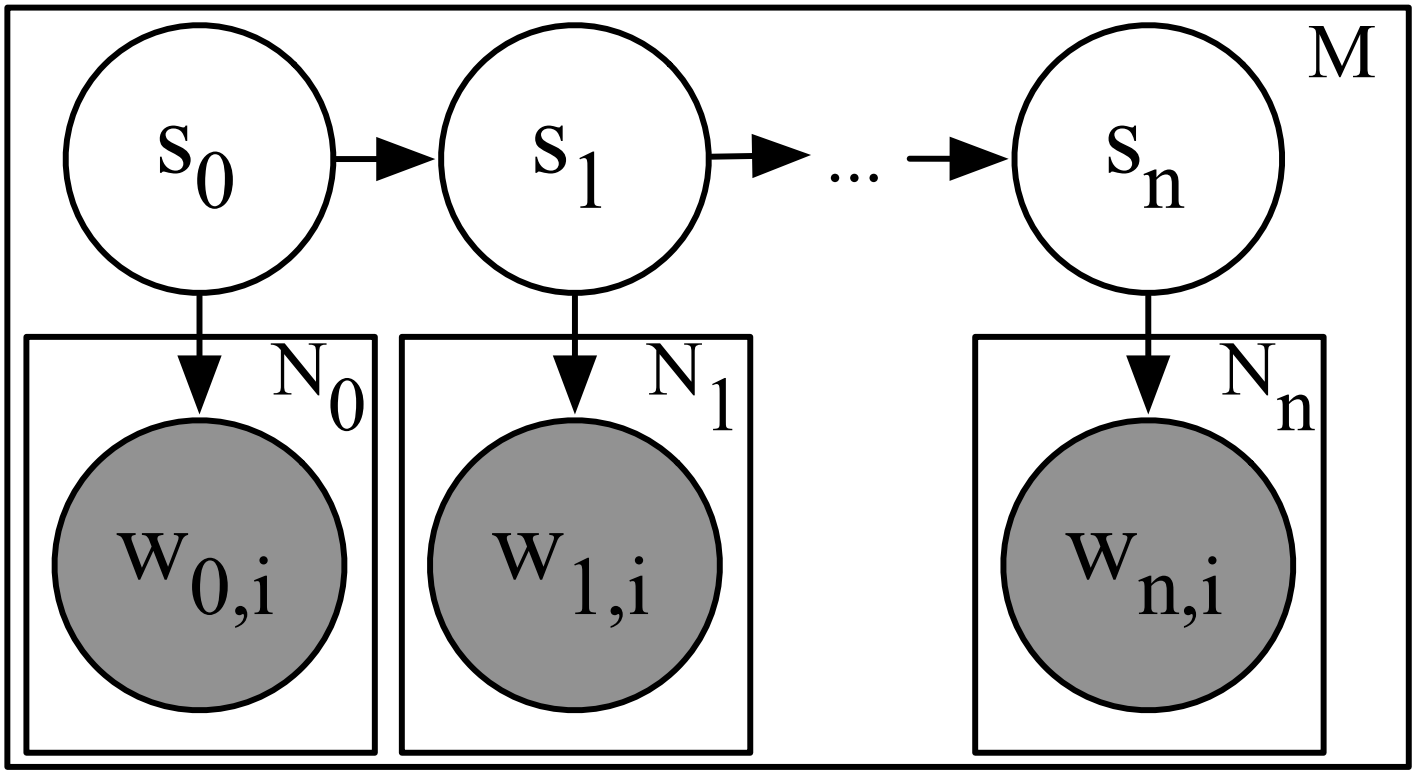

TM-HMM: Share Topics across States

~~LM~~TM-HMM: HMM with state ~~language~~ topic model

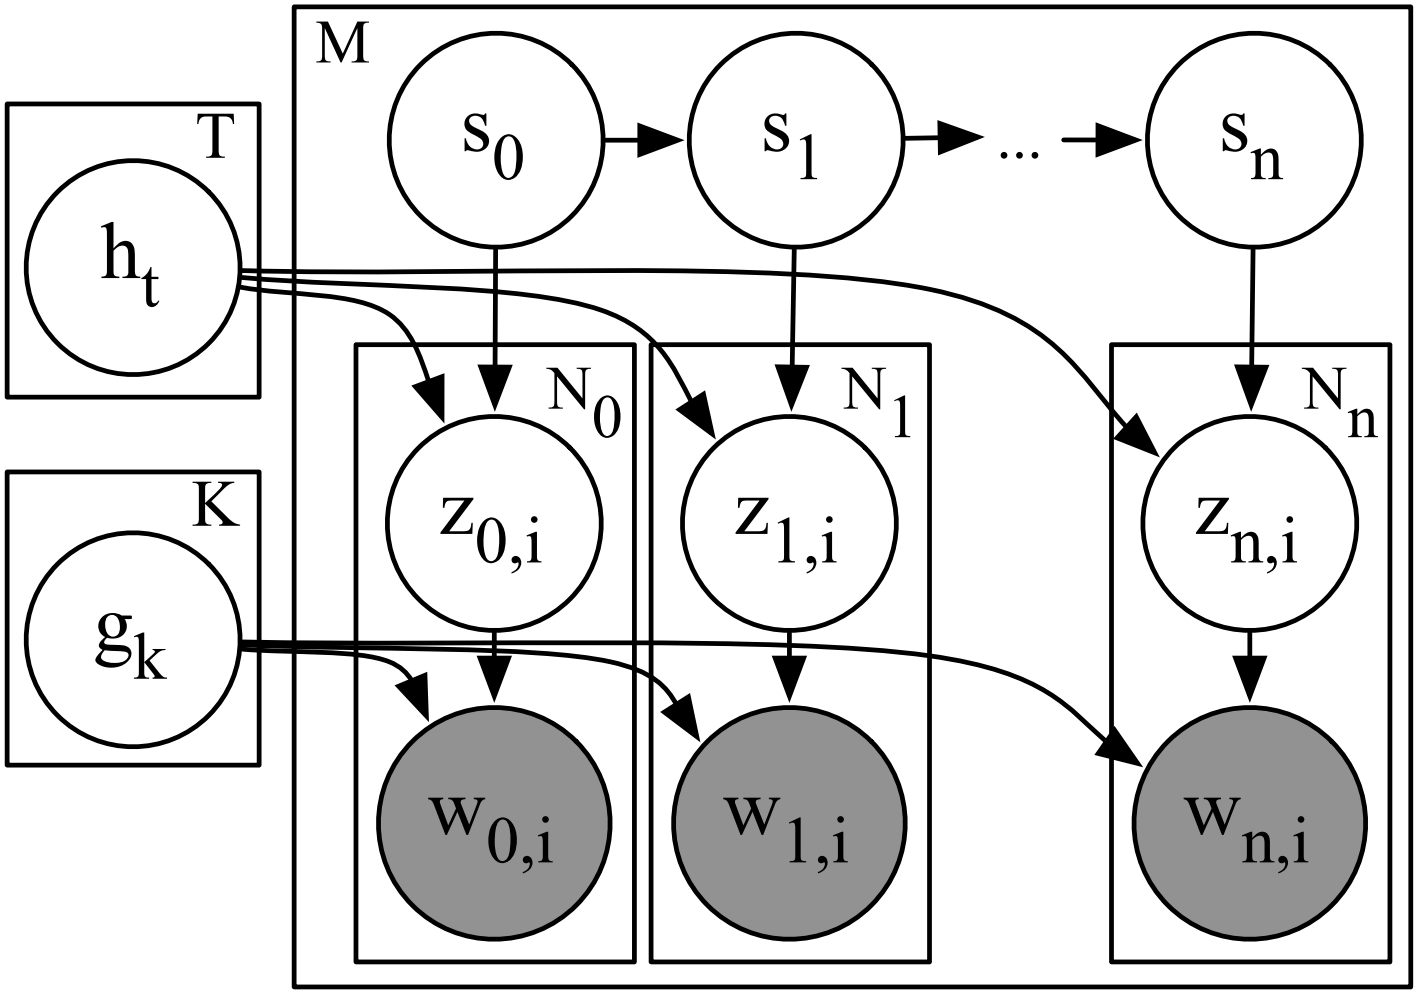

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

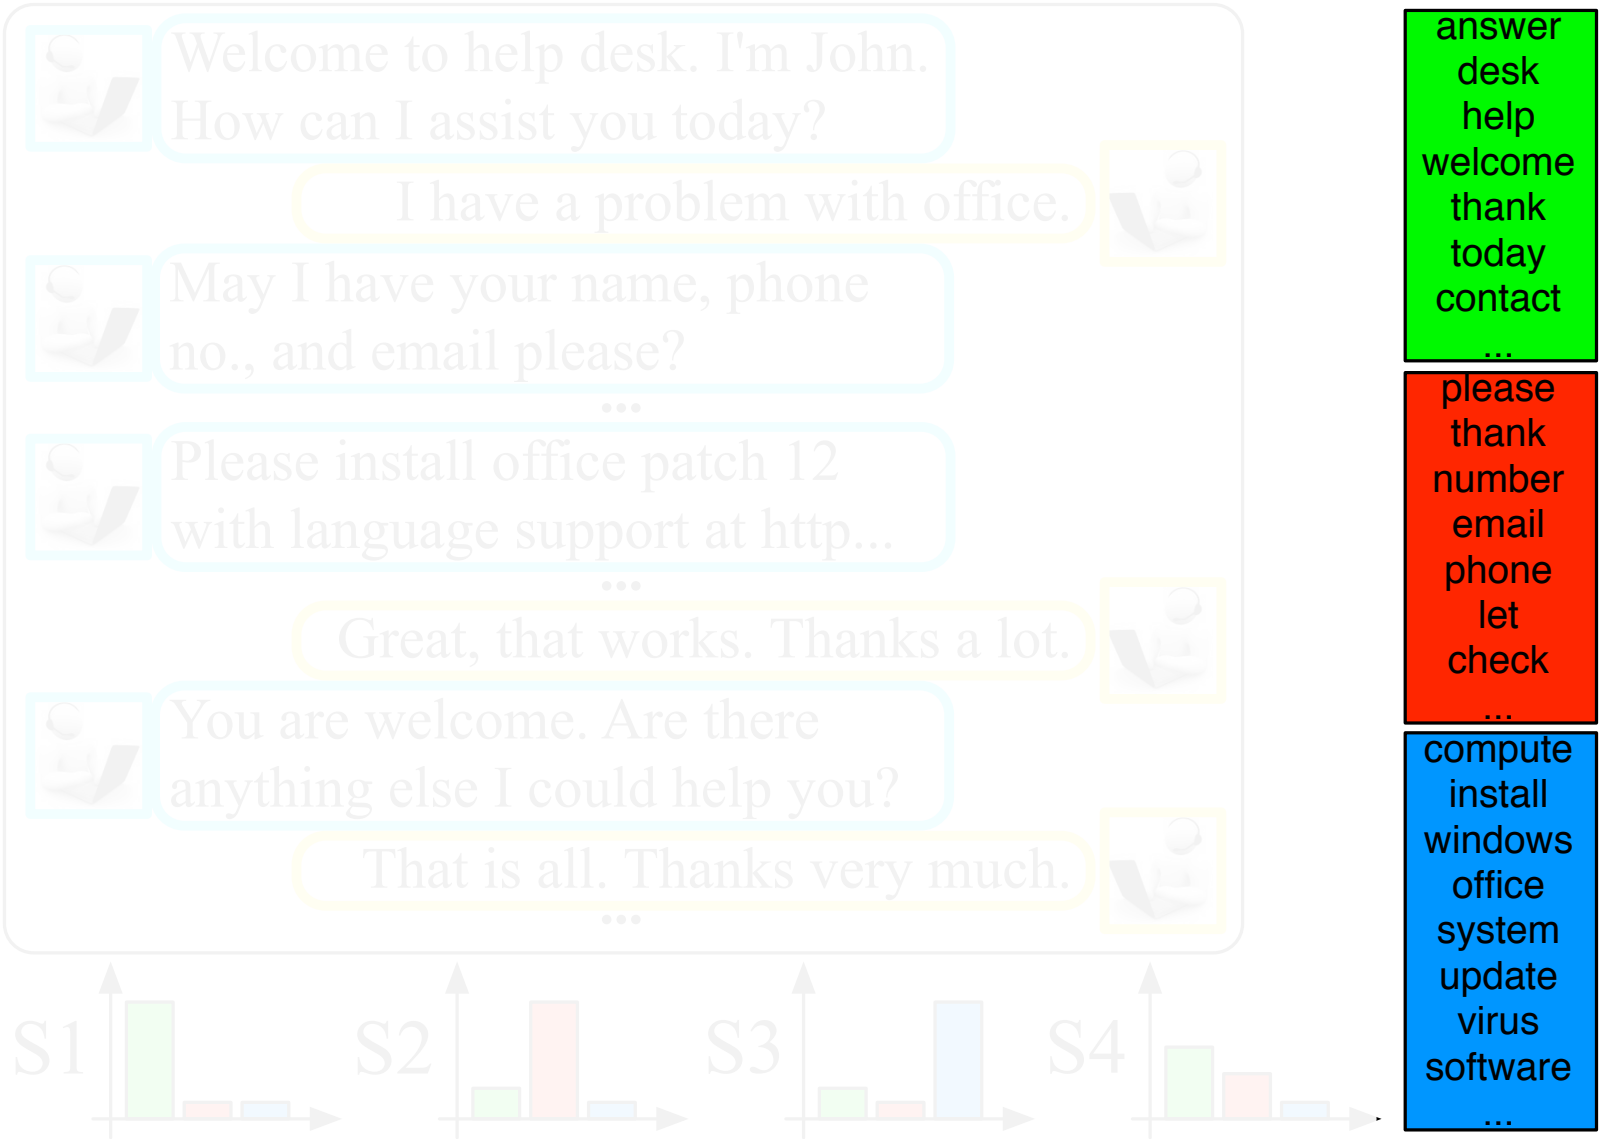

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

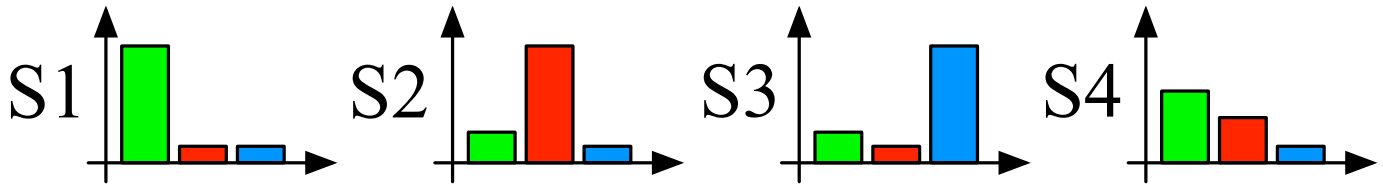

answer
desk
help
welcome
thank
today
contact
...

please
thank
number
email
phone
let
check
...

compute
install
windows
office
system
update
virus
software
...

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

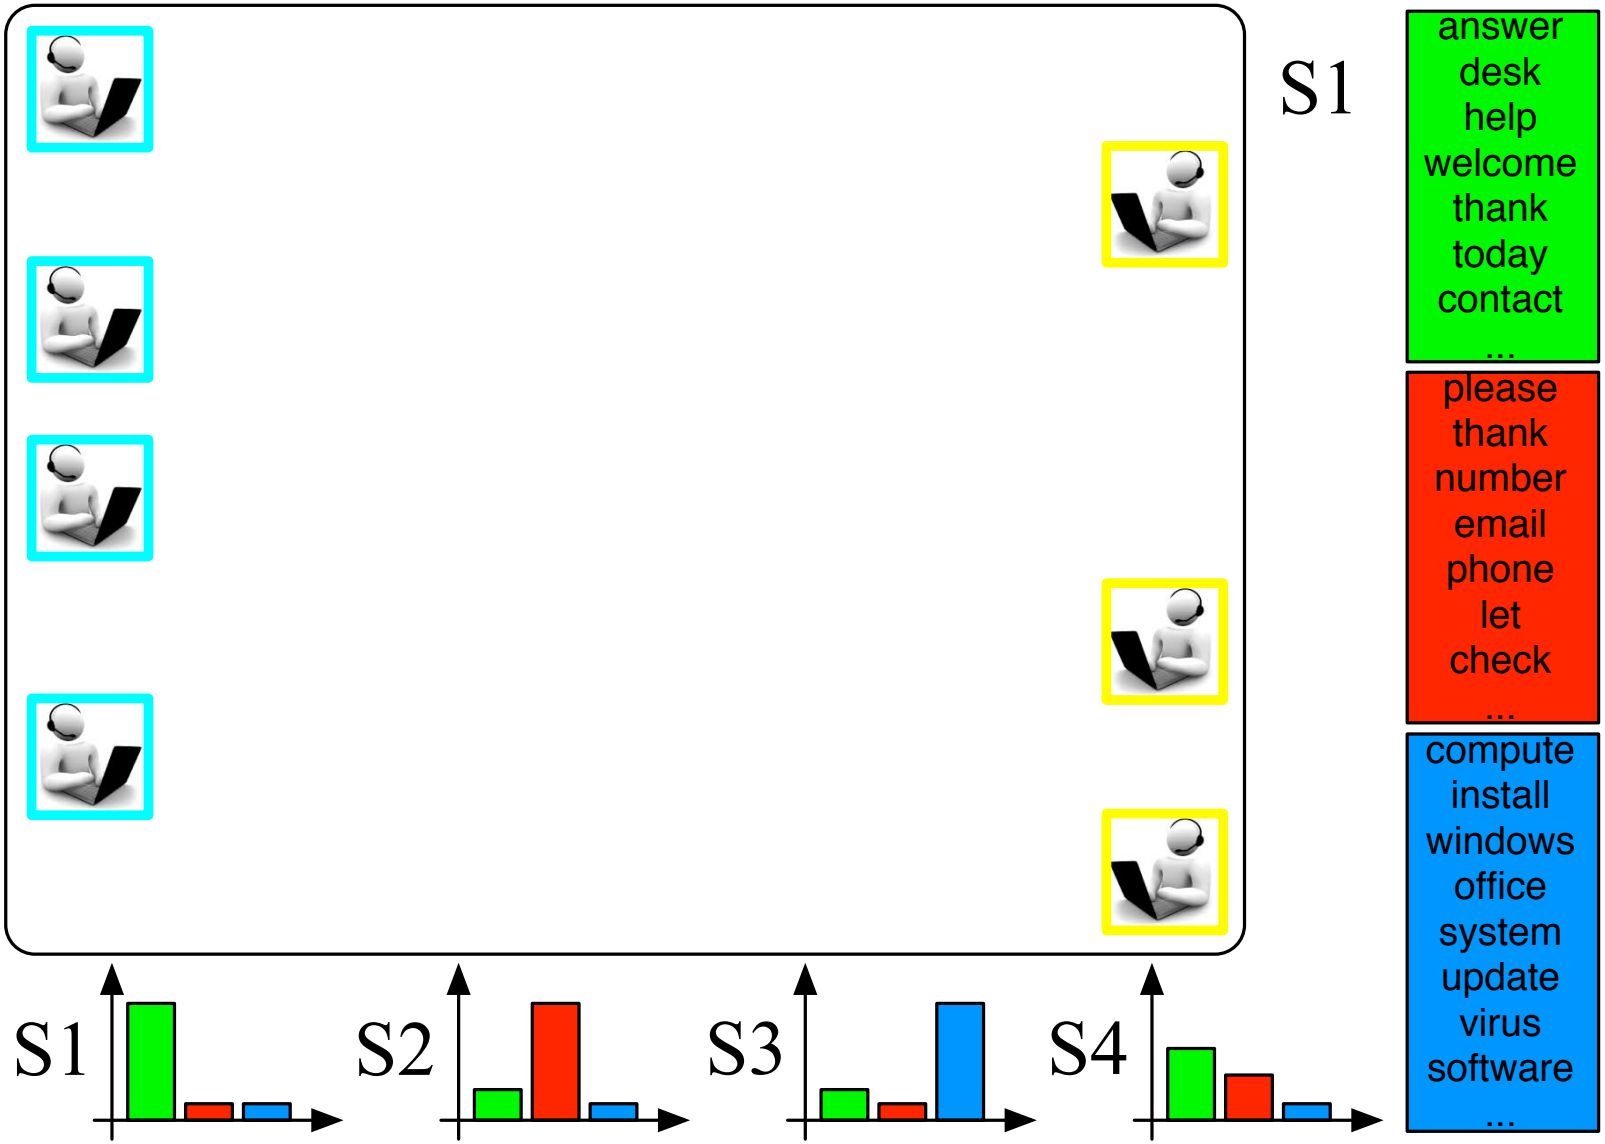

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

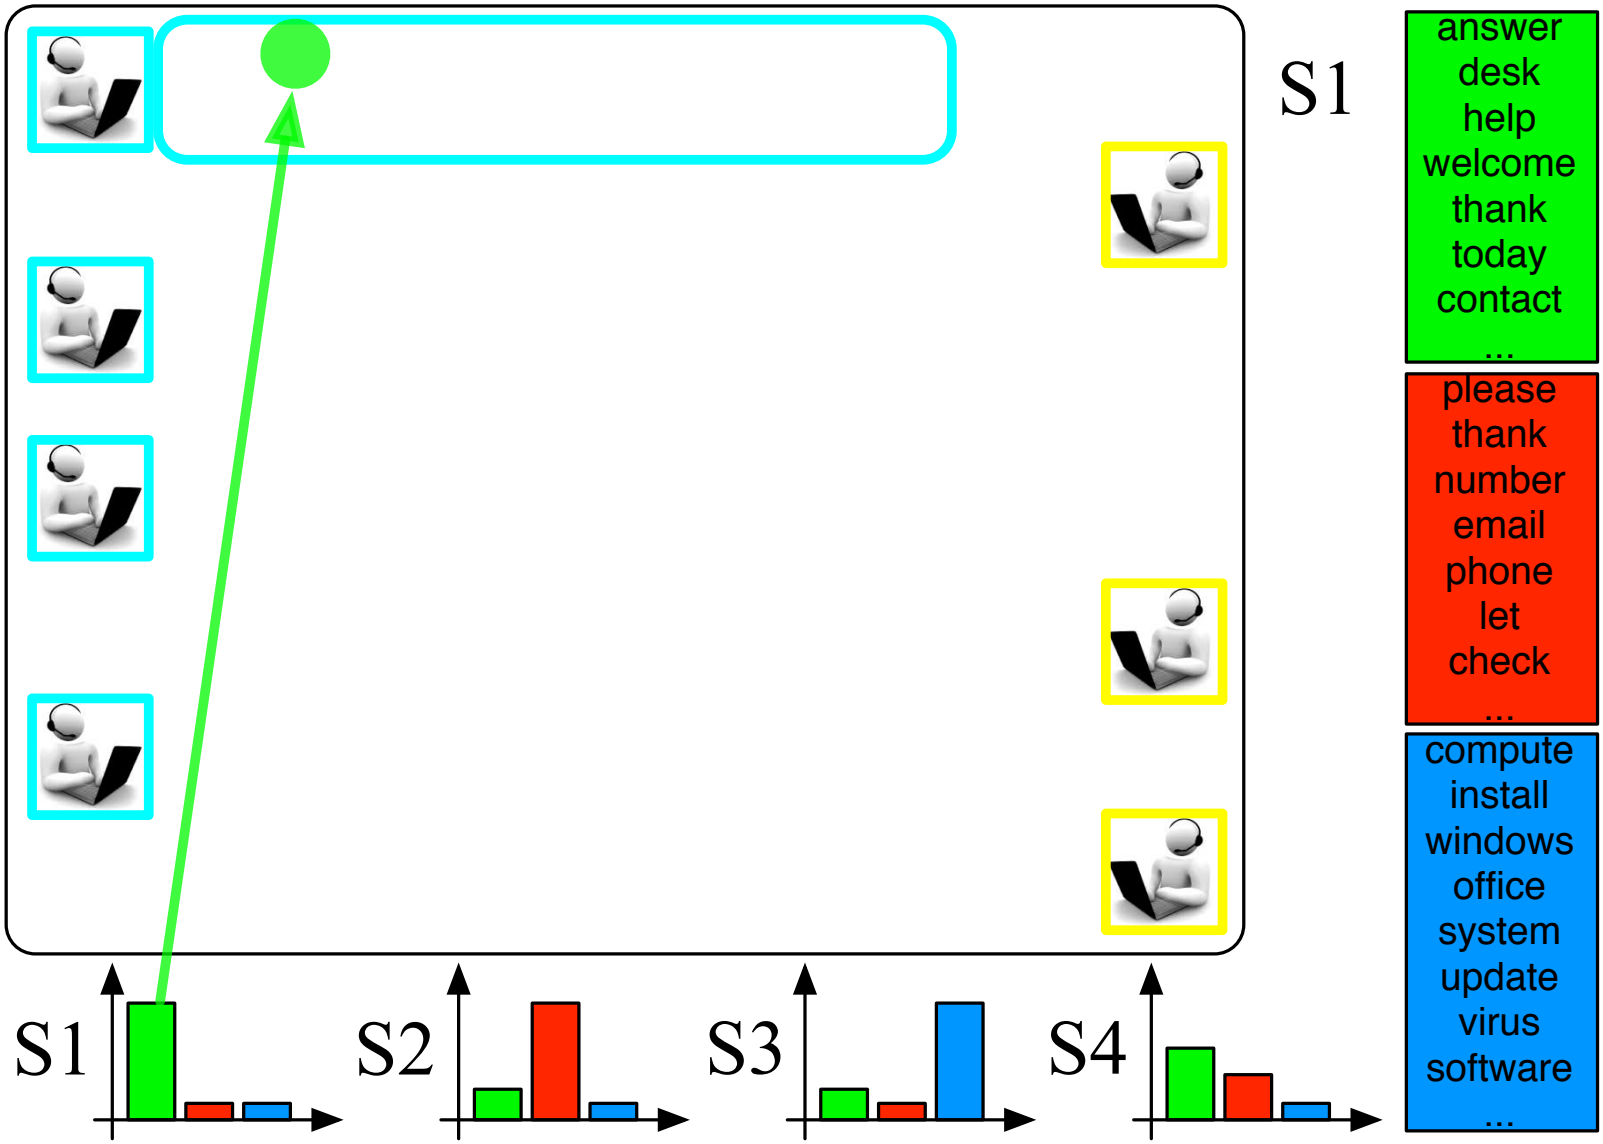

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

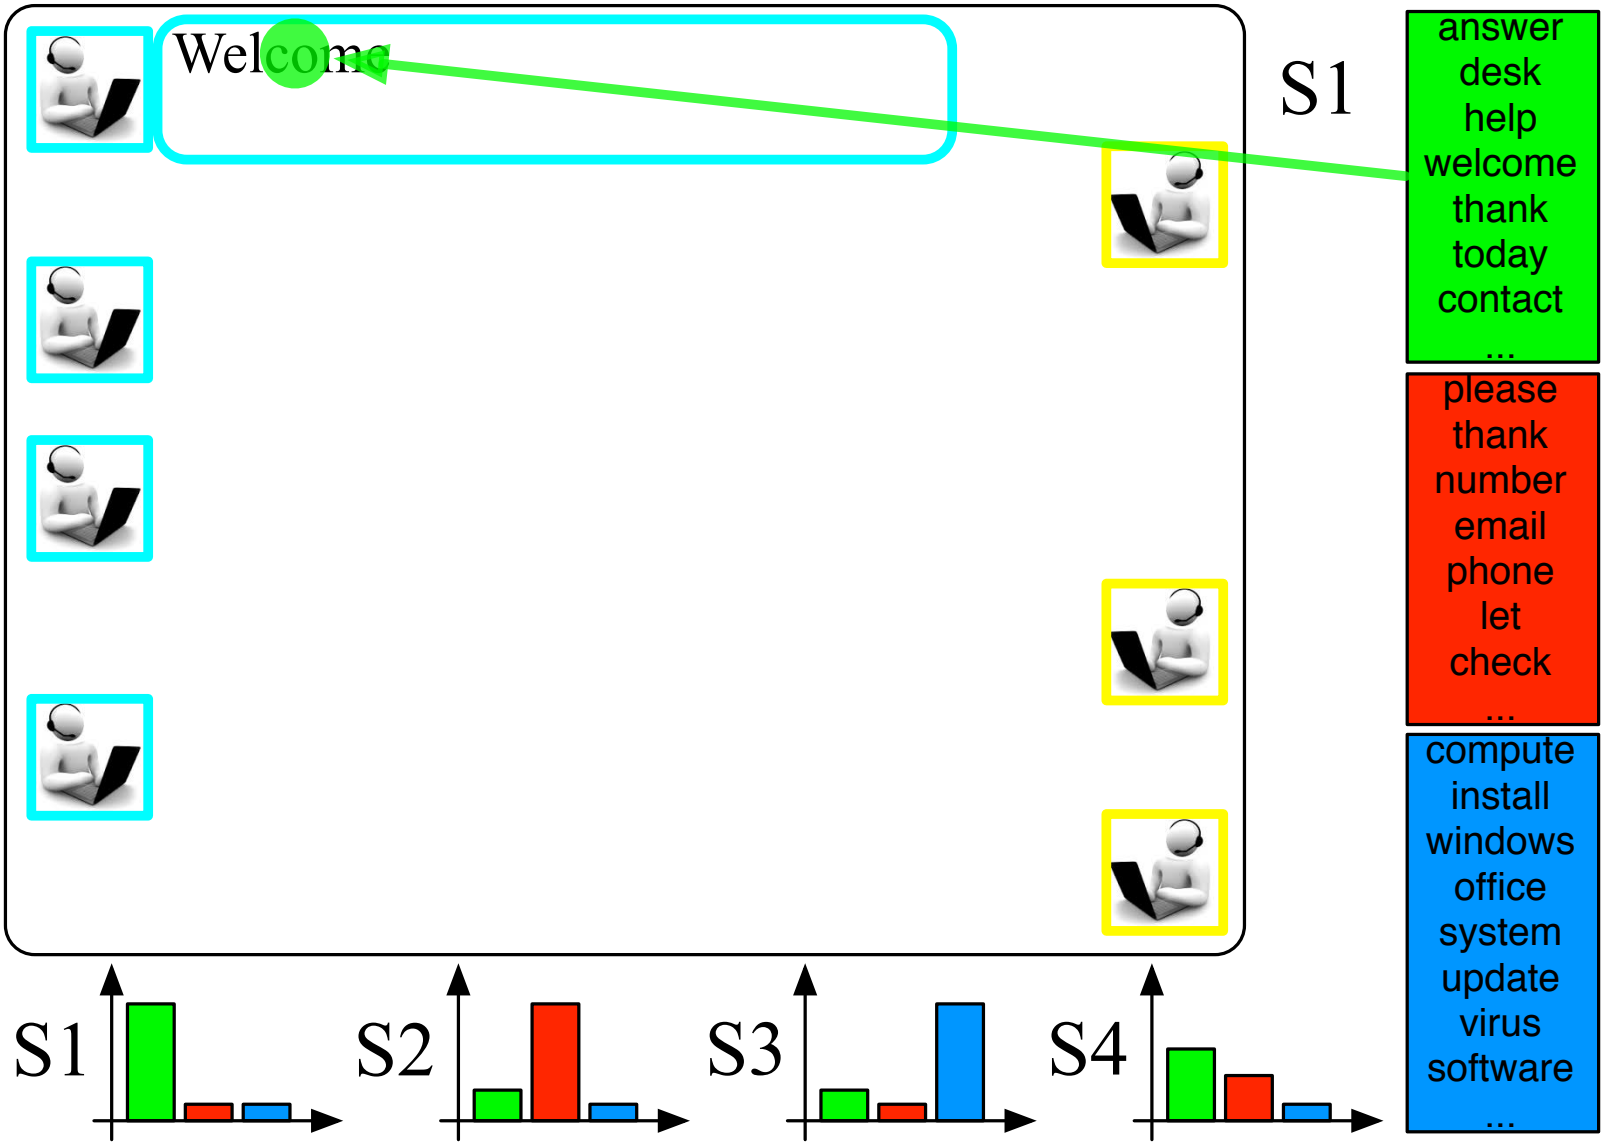

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

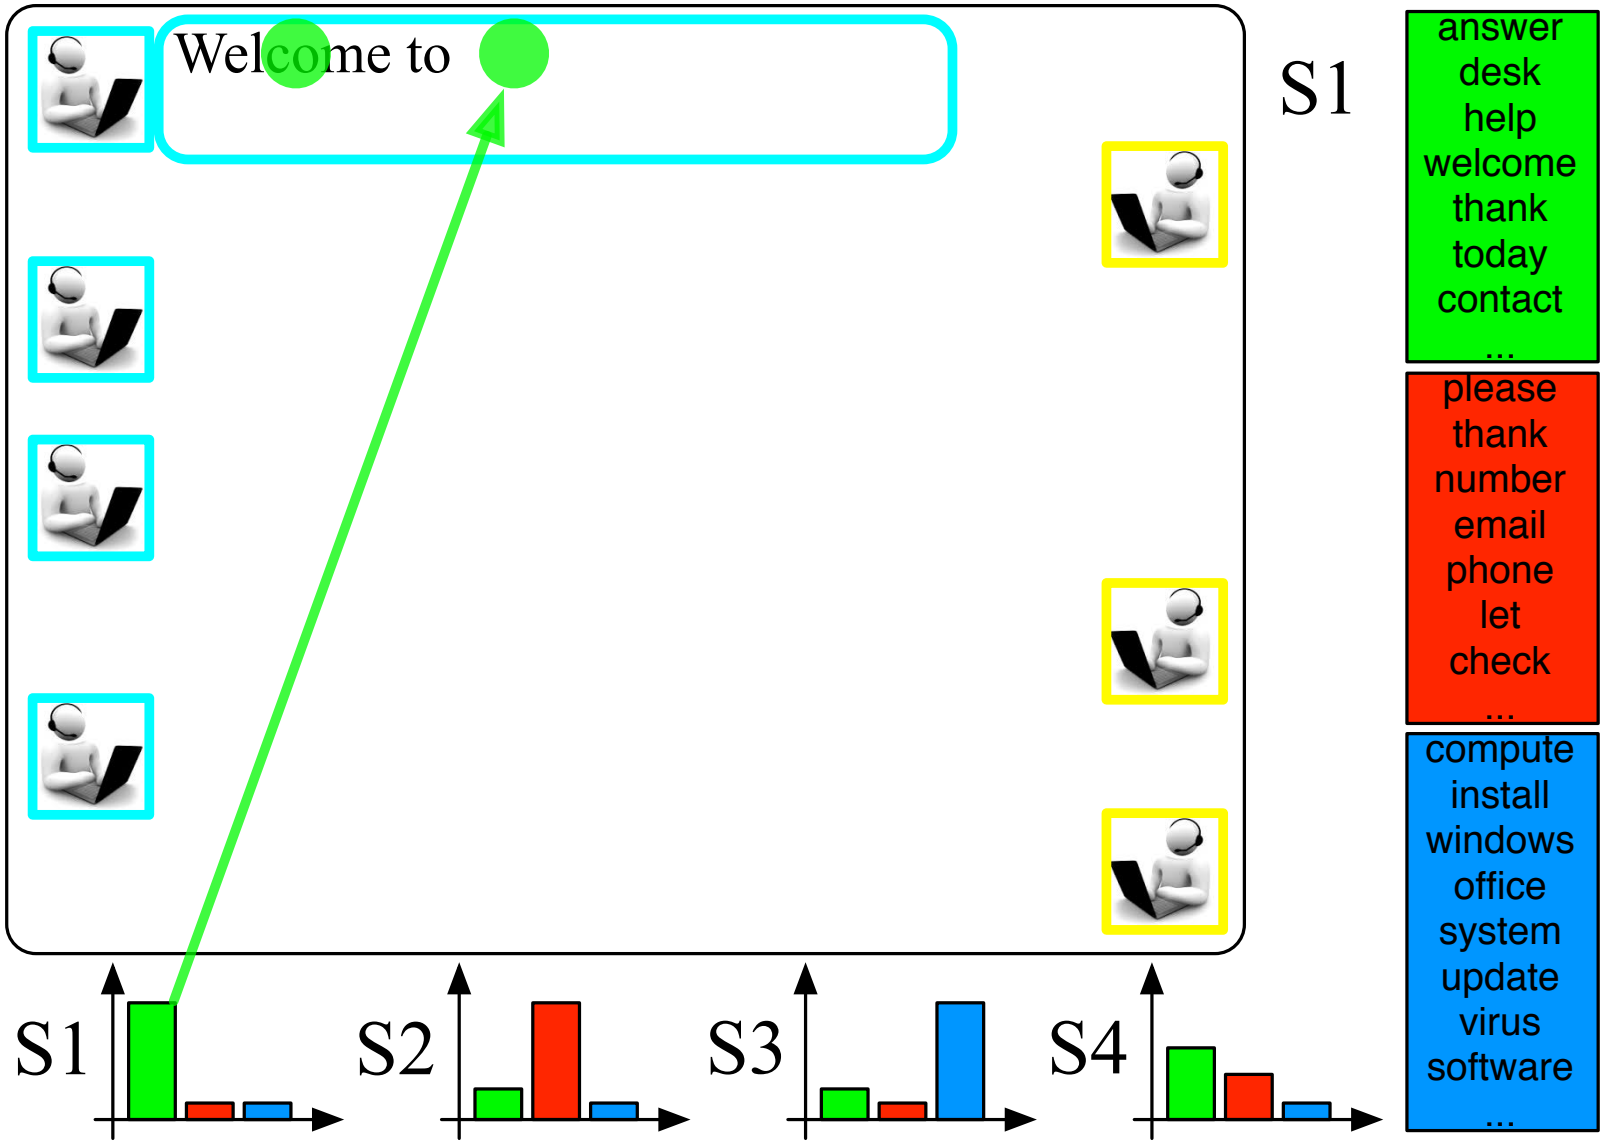

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

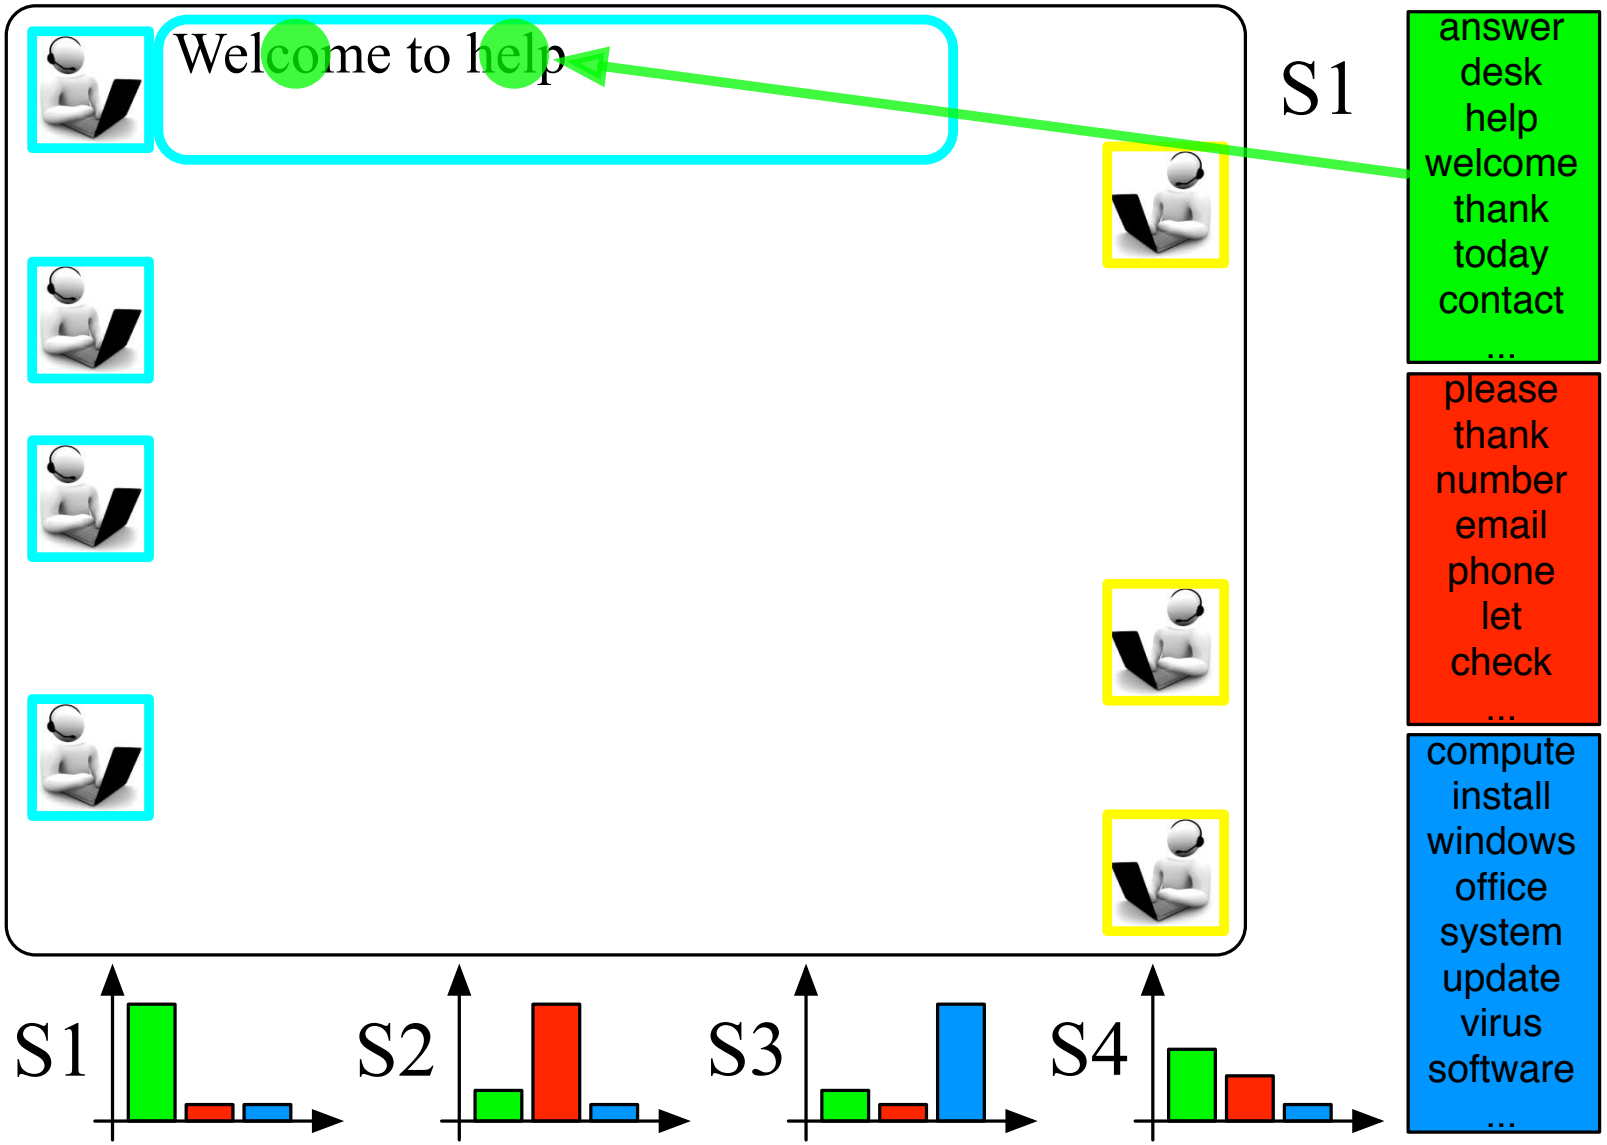

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

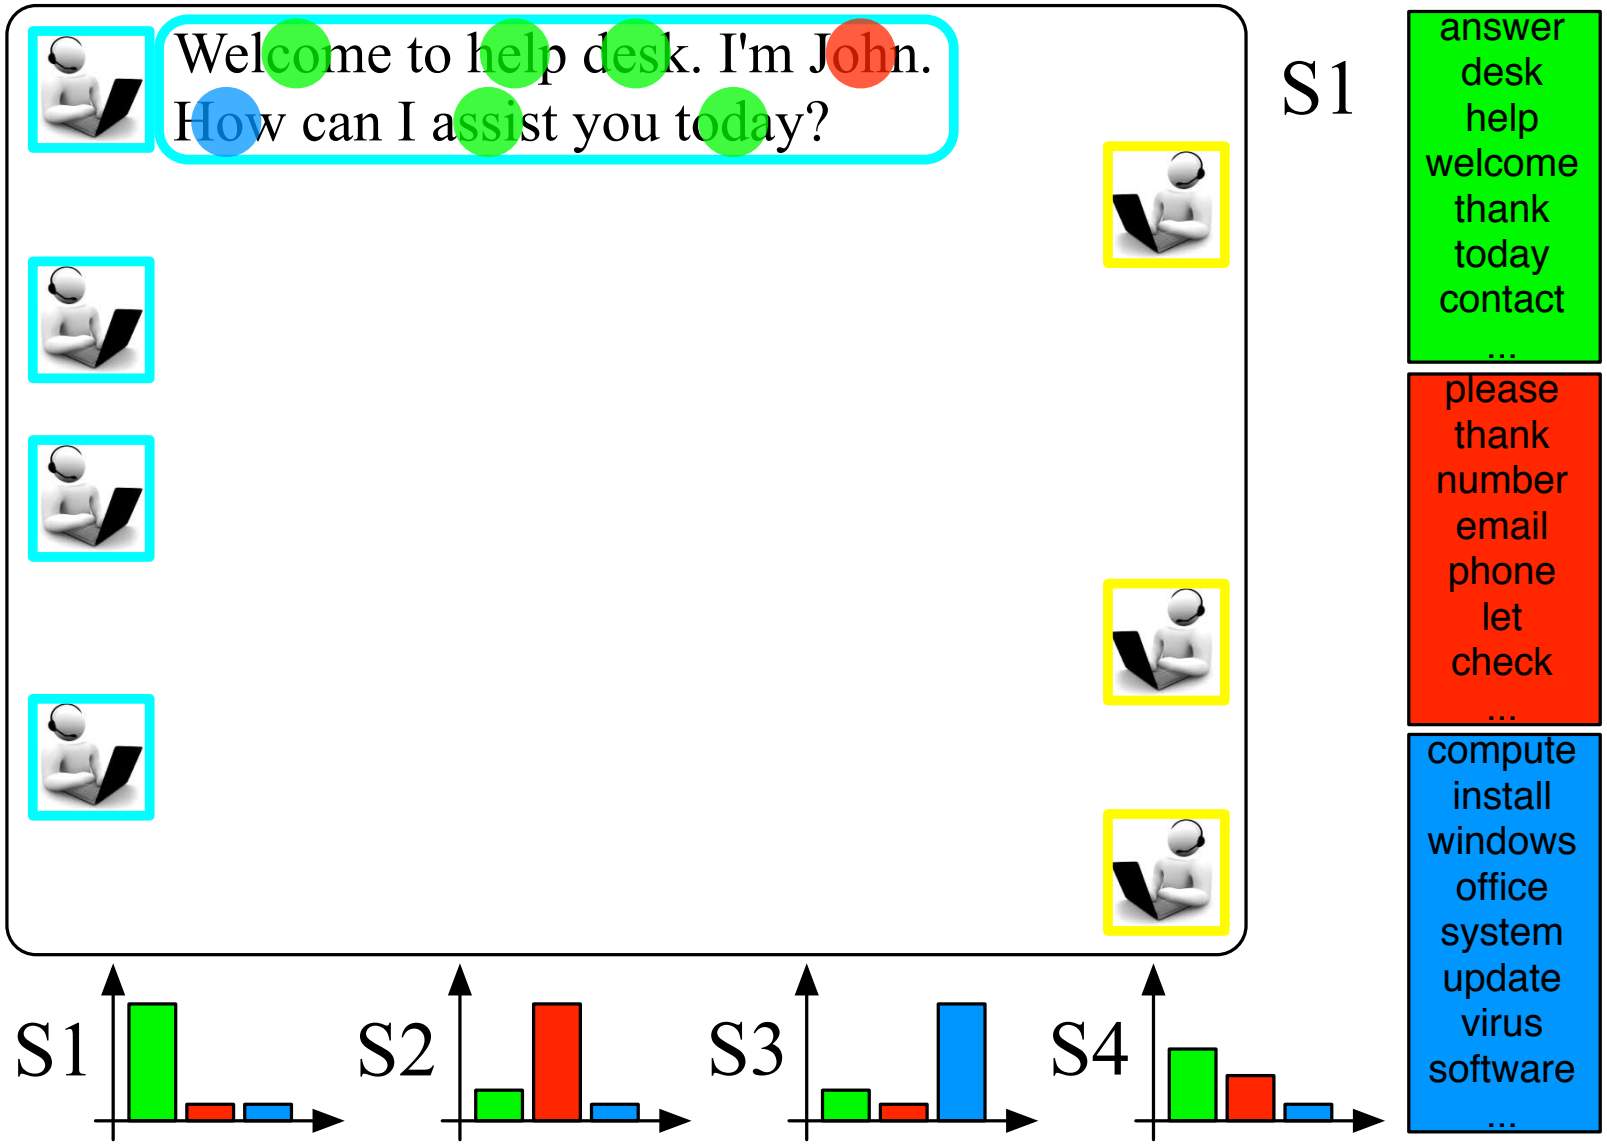

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

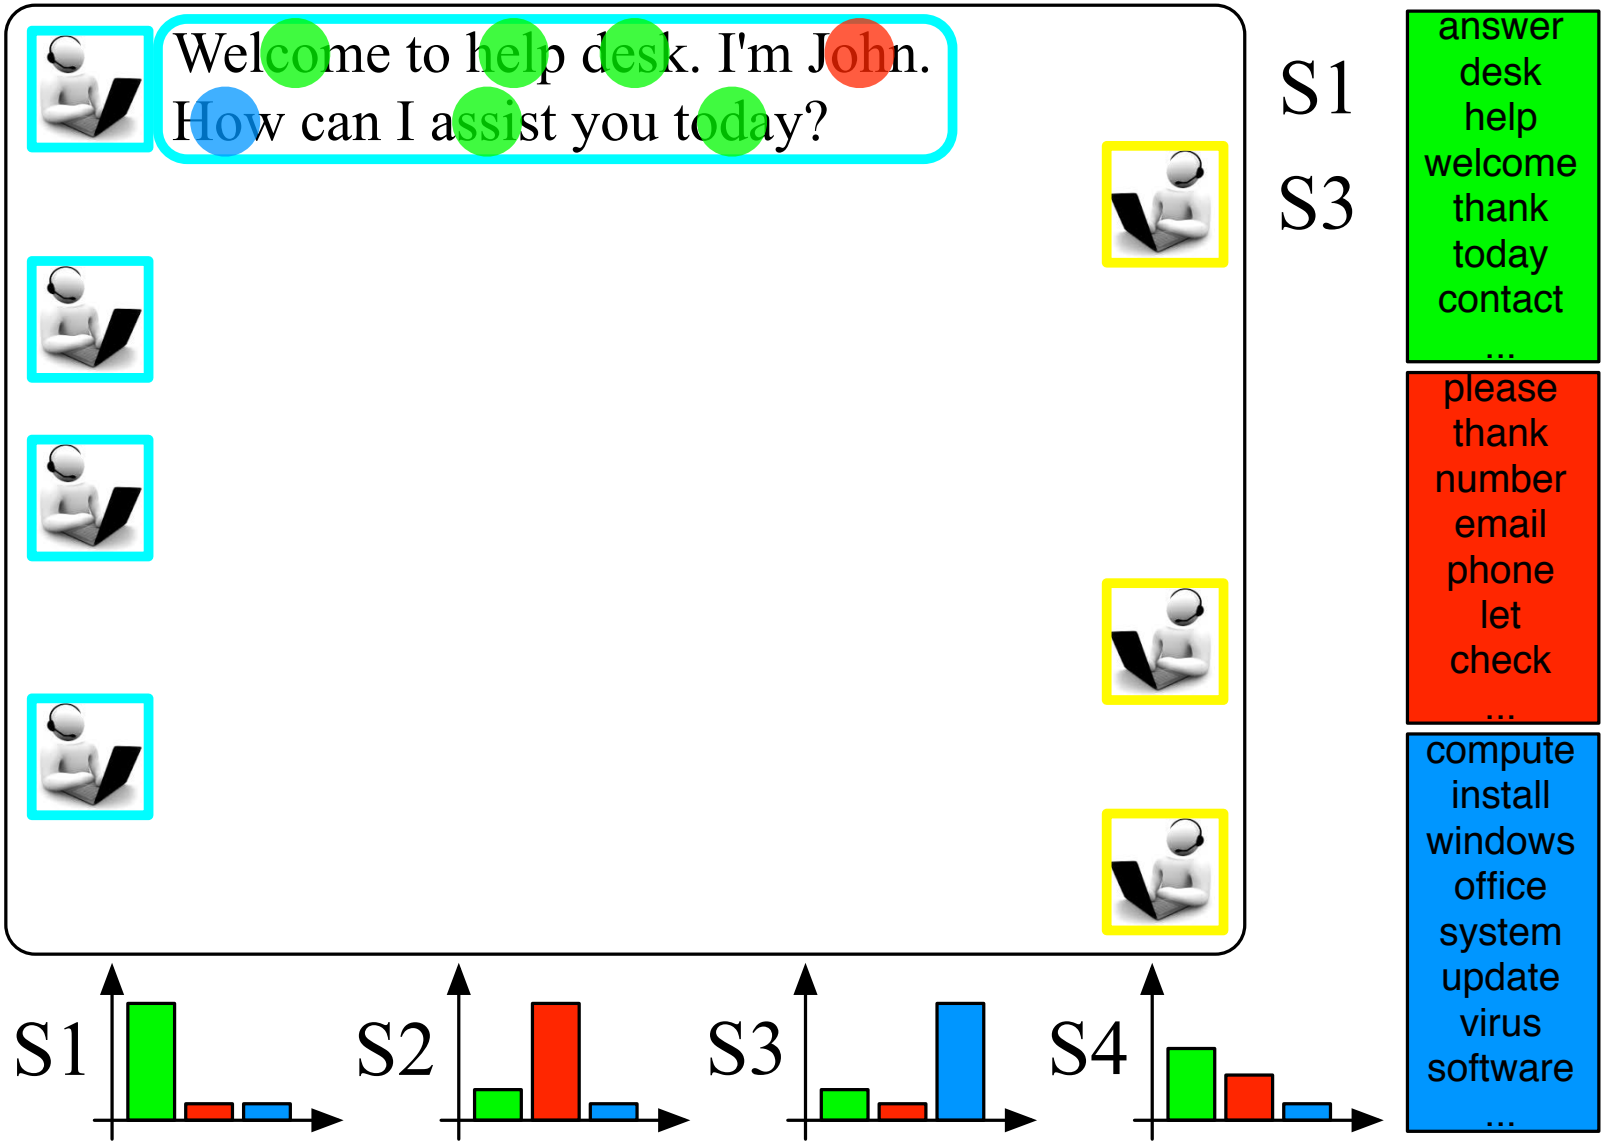

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

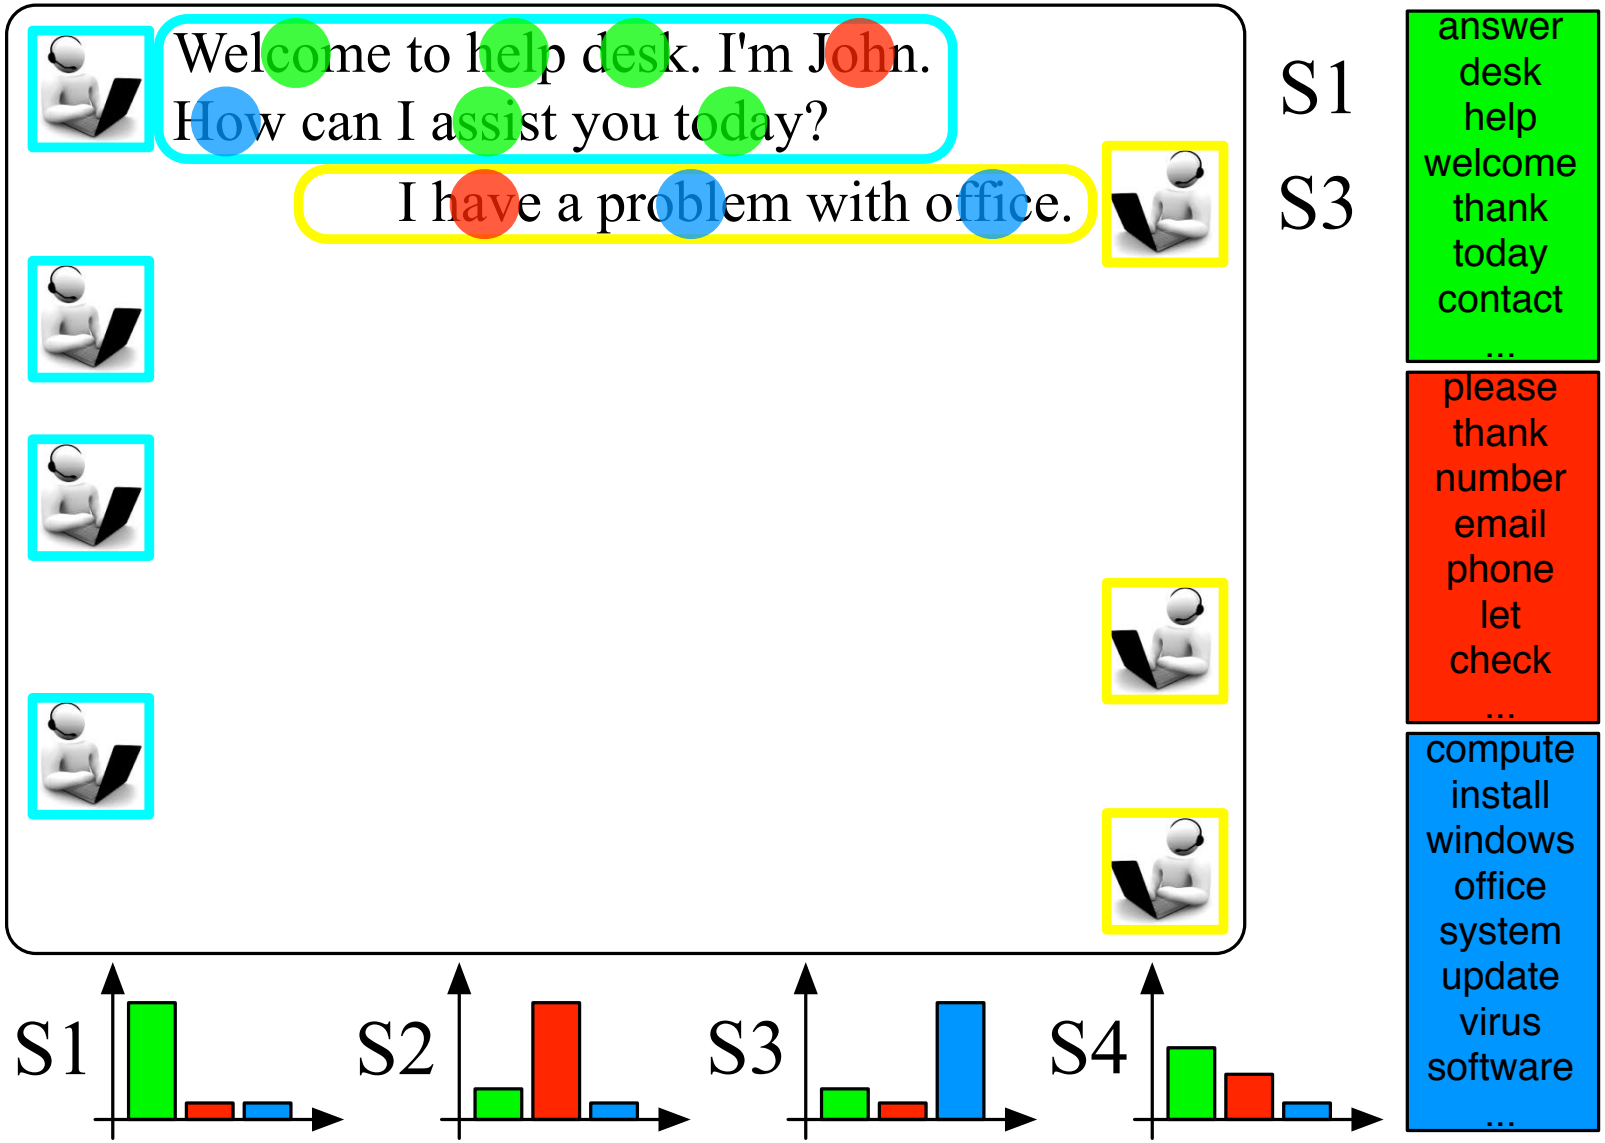

TM-HMM: Share Topics across States

TM-HMM: HMM with state topic model

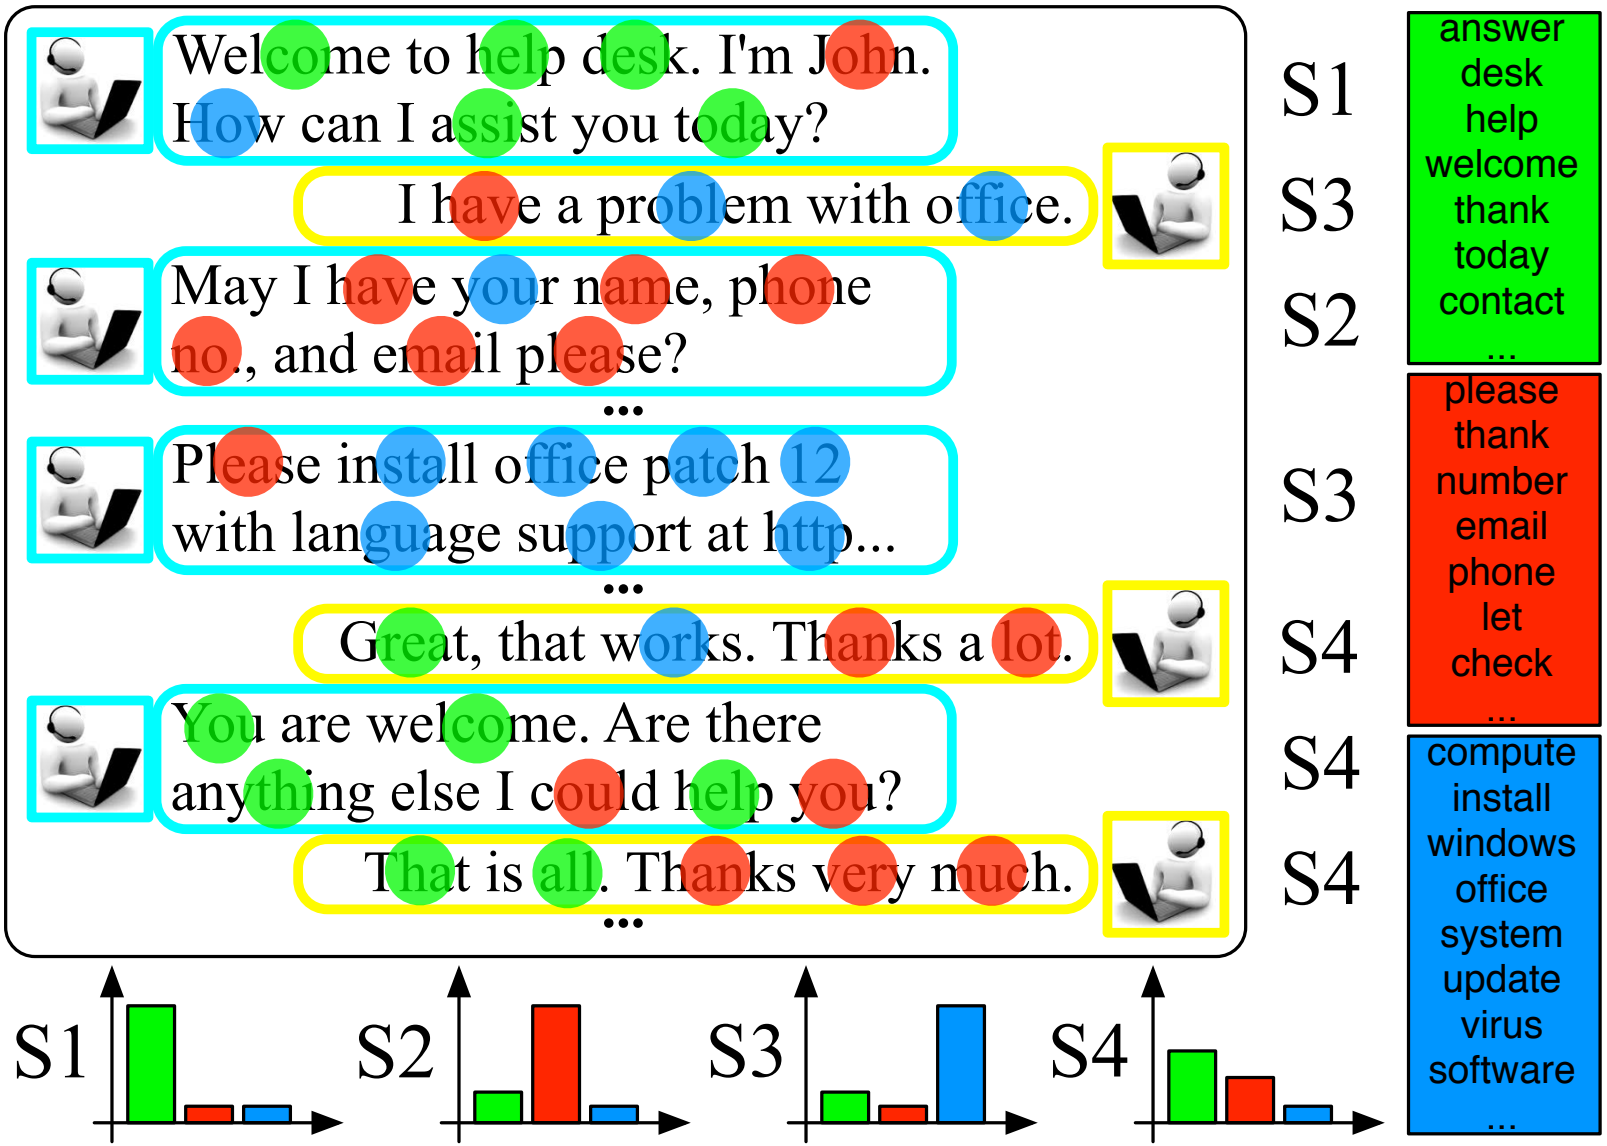

TM-HMMSS: Share Topics across Sessions

LM-HMMS: HMM with state language model and source generator

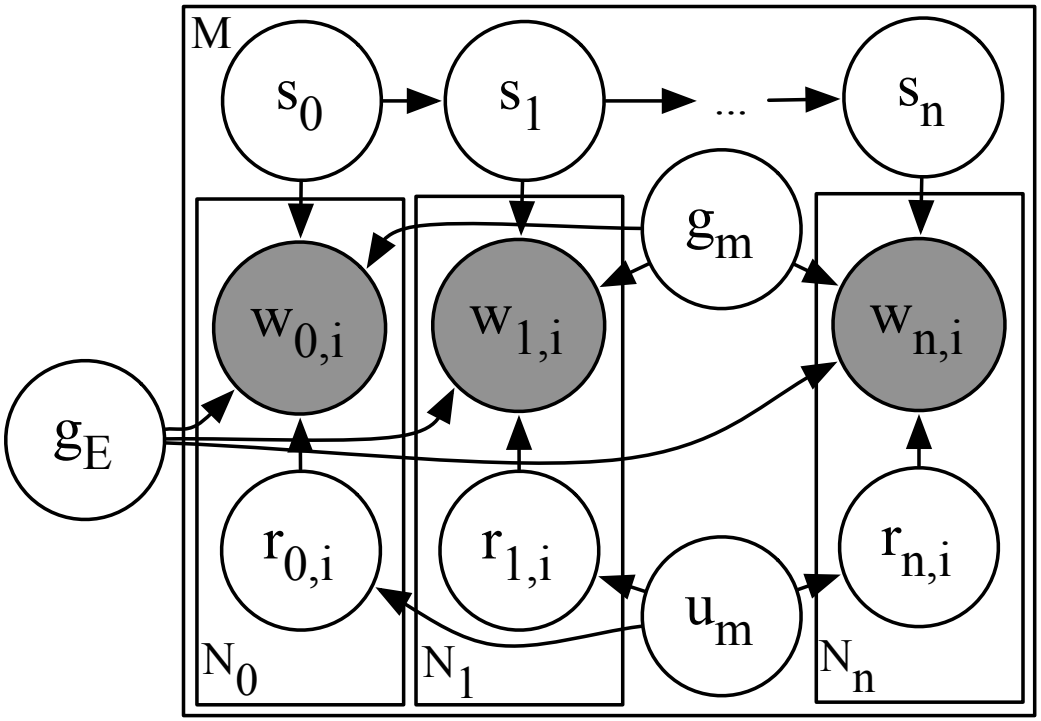

TM-HMMSS: Share Topics across Sessions

~~LM~~TM-HMMSS: HMM with state language model
and state source generator and session topic model

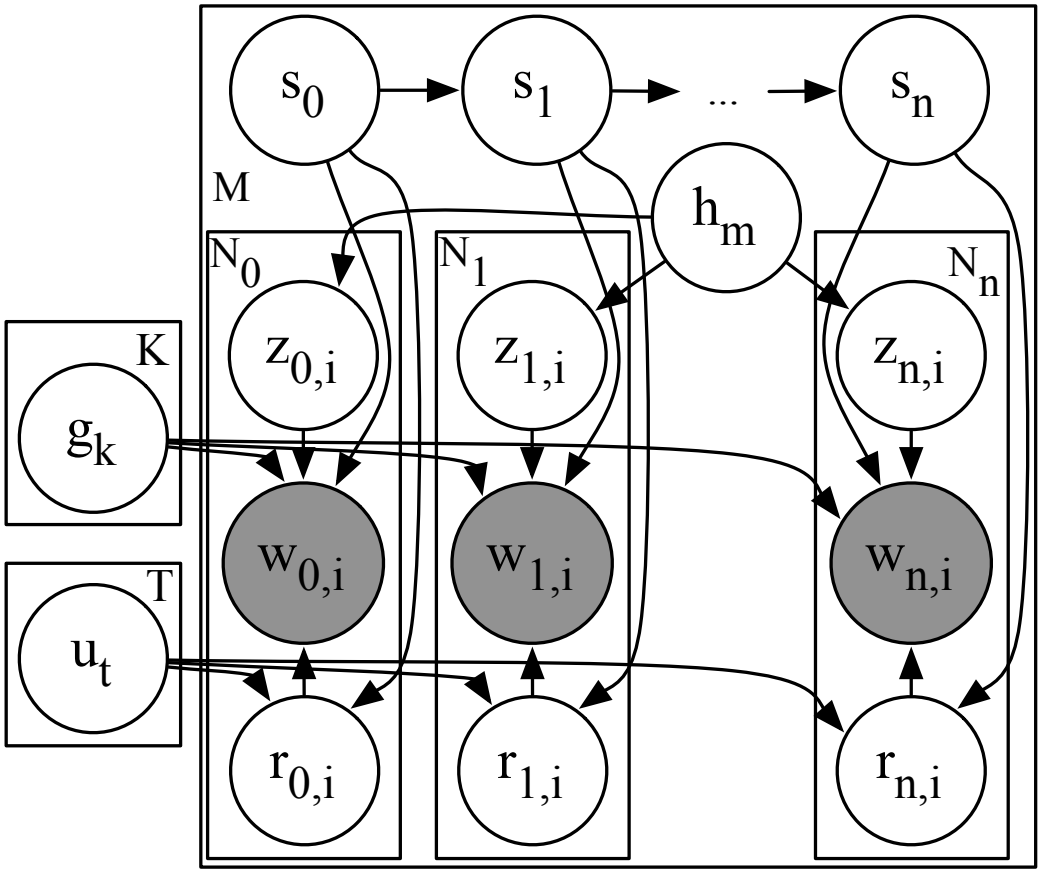

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

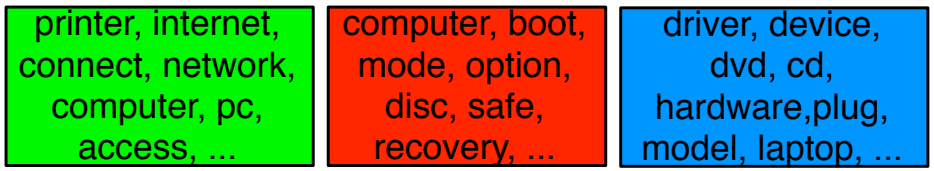

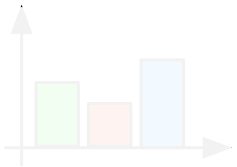

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

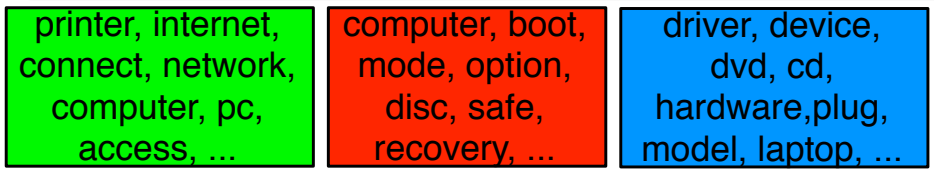

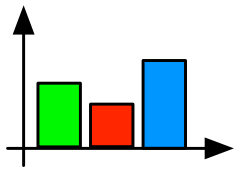

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model
and state source generator and session topic model

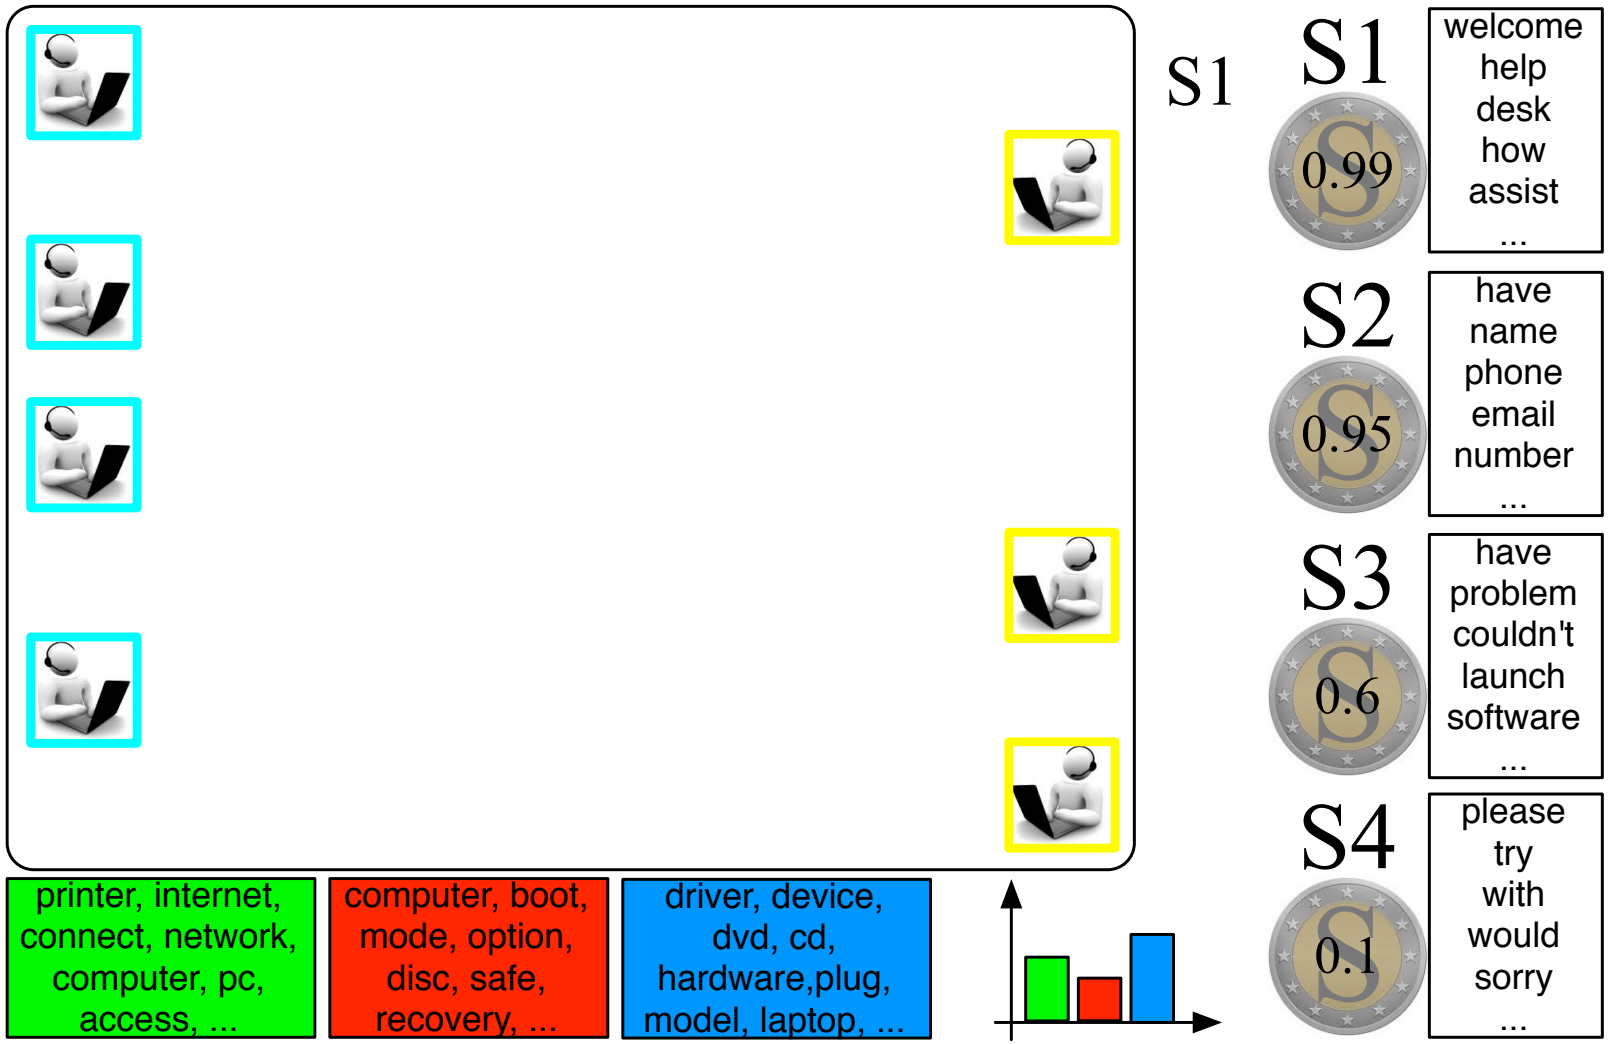

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

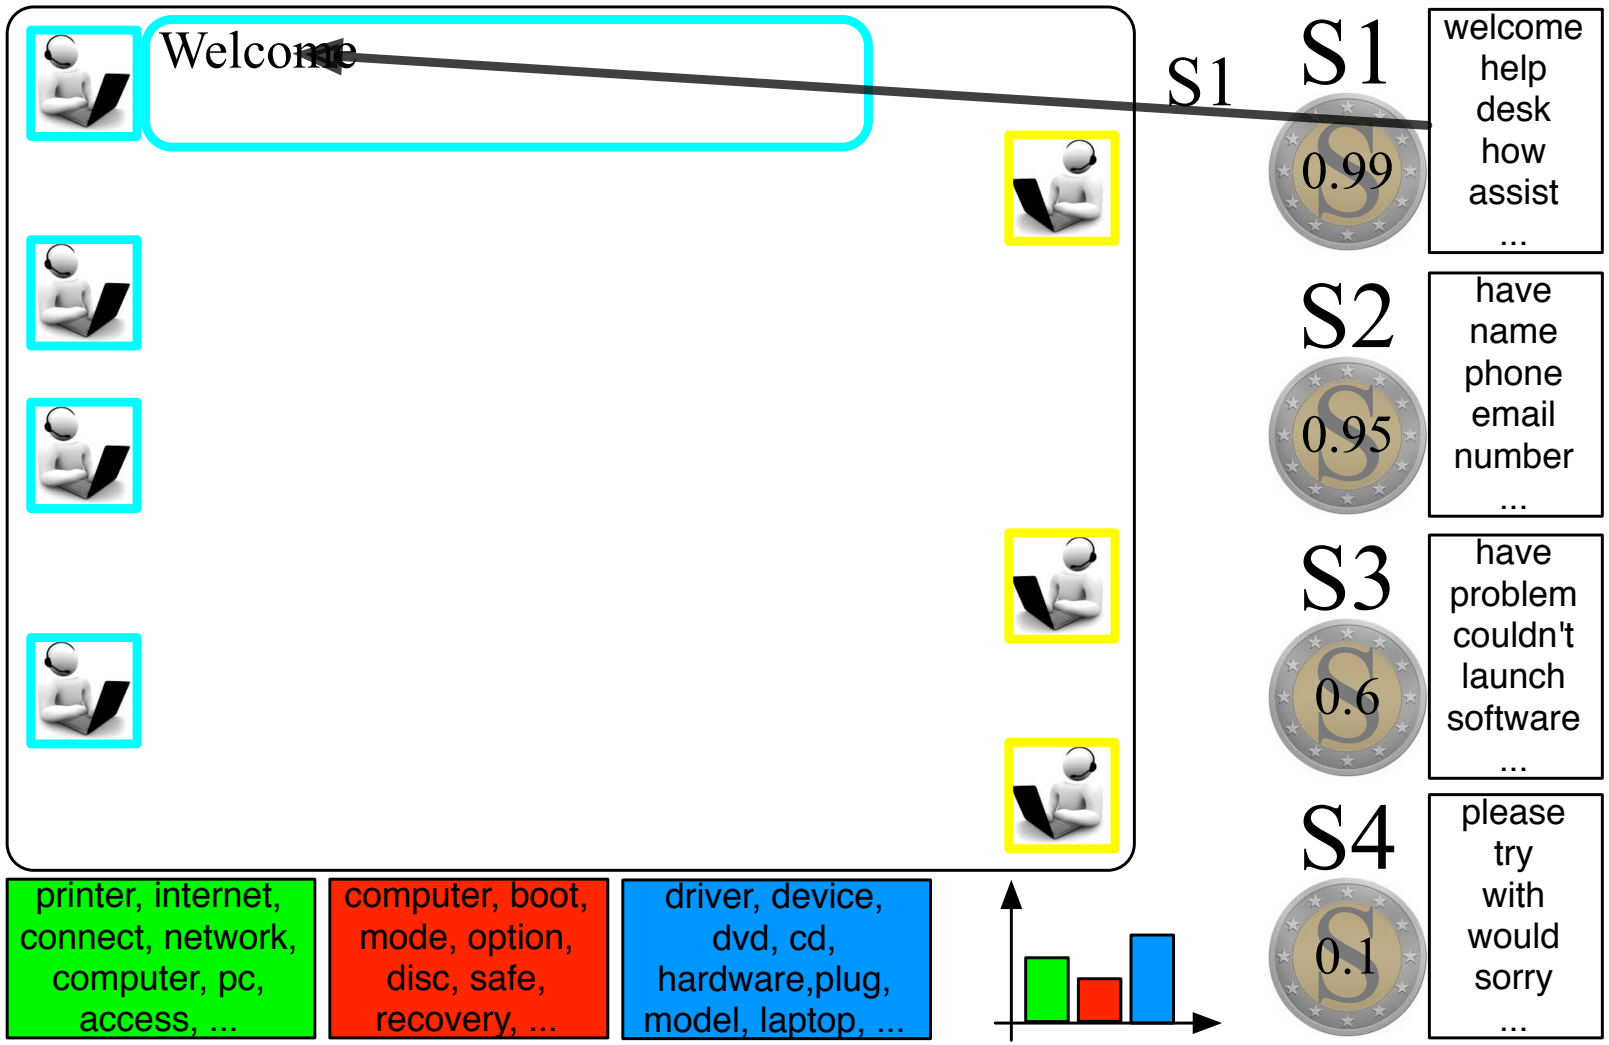

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

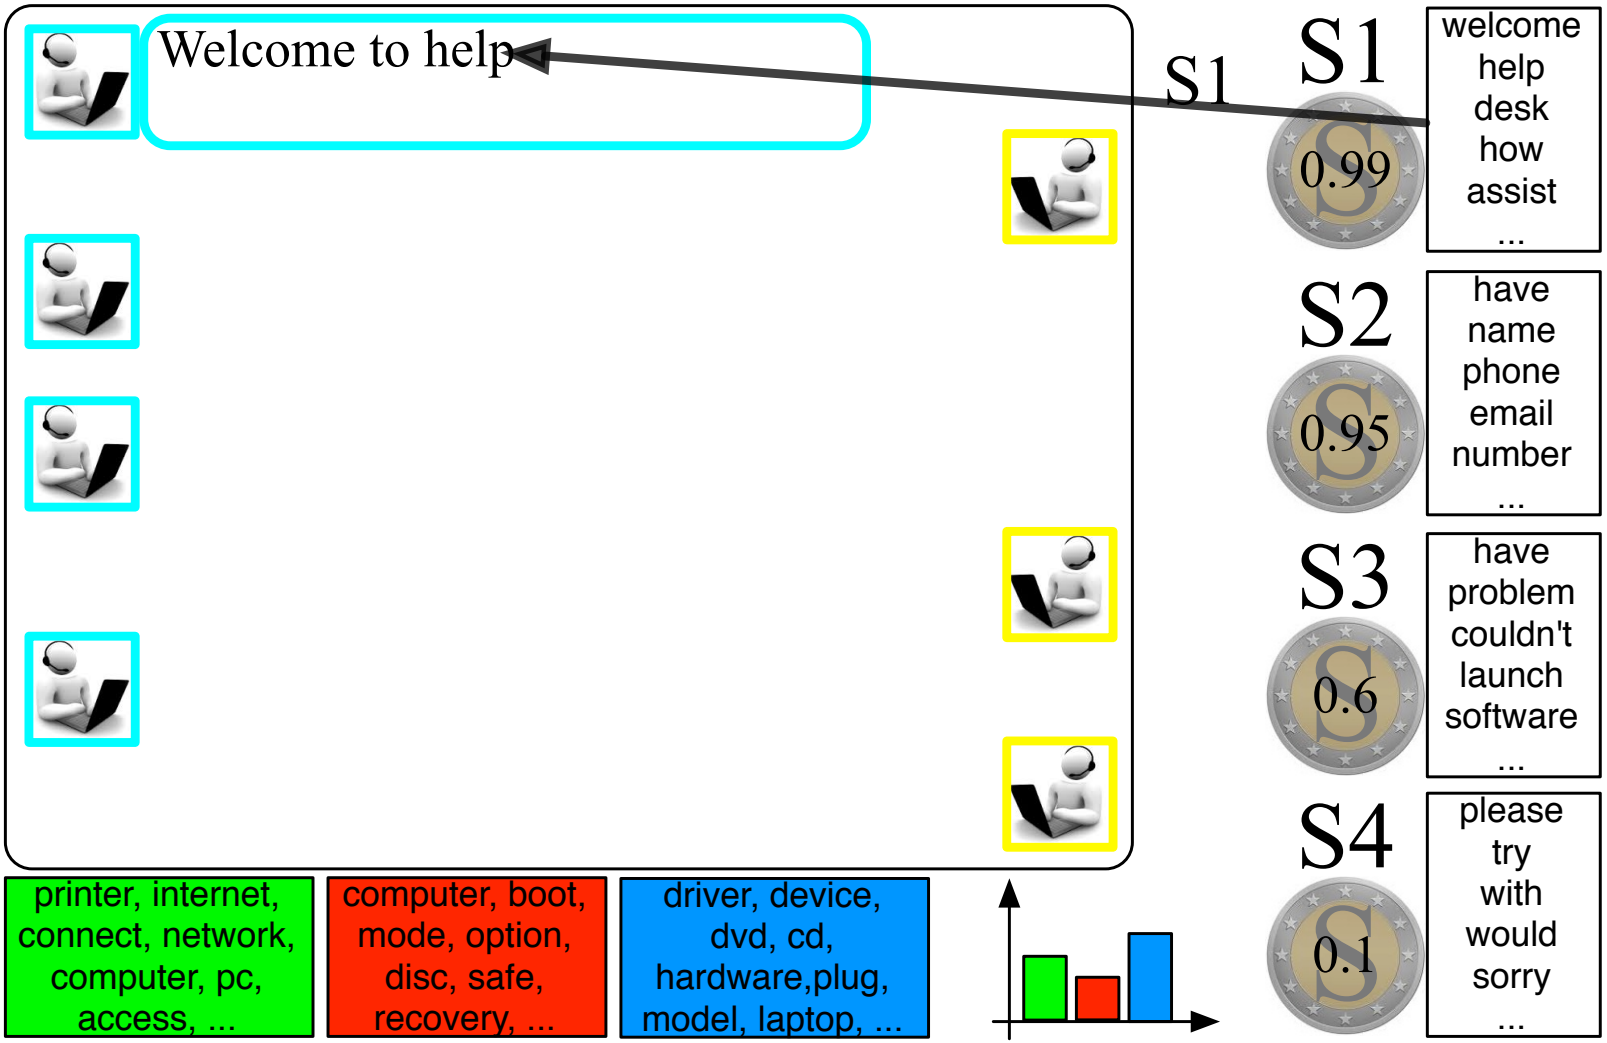

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

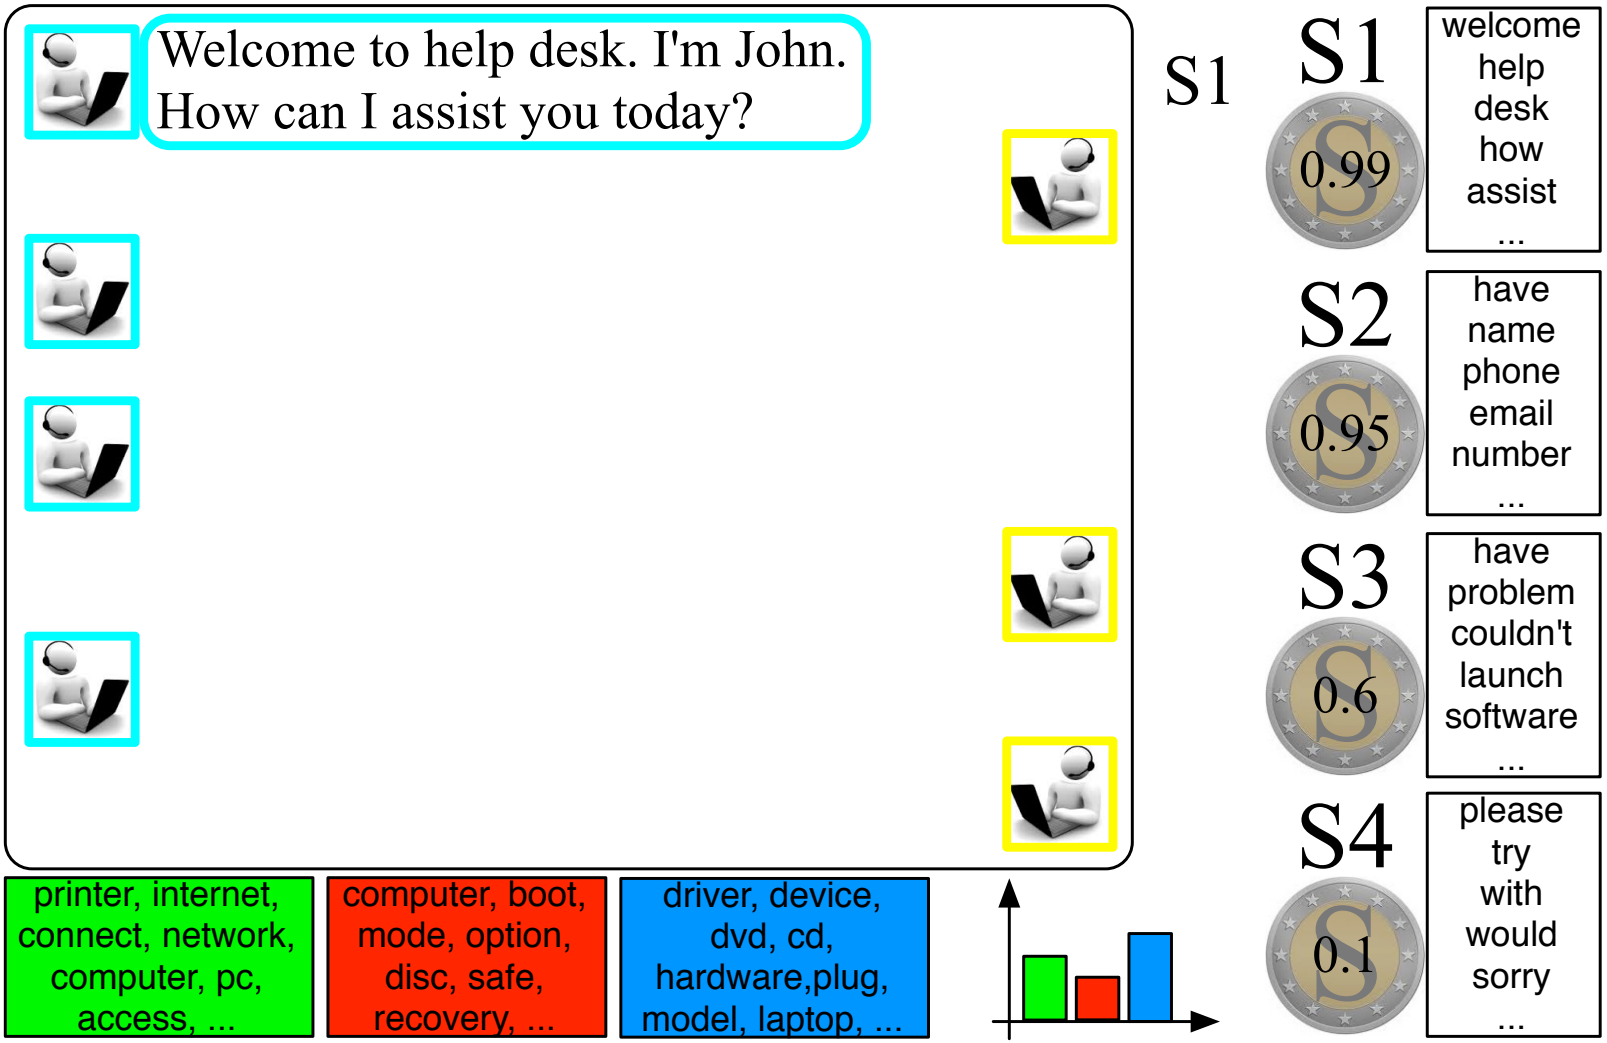

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model
and state source generator and session topic model

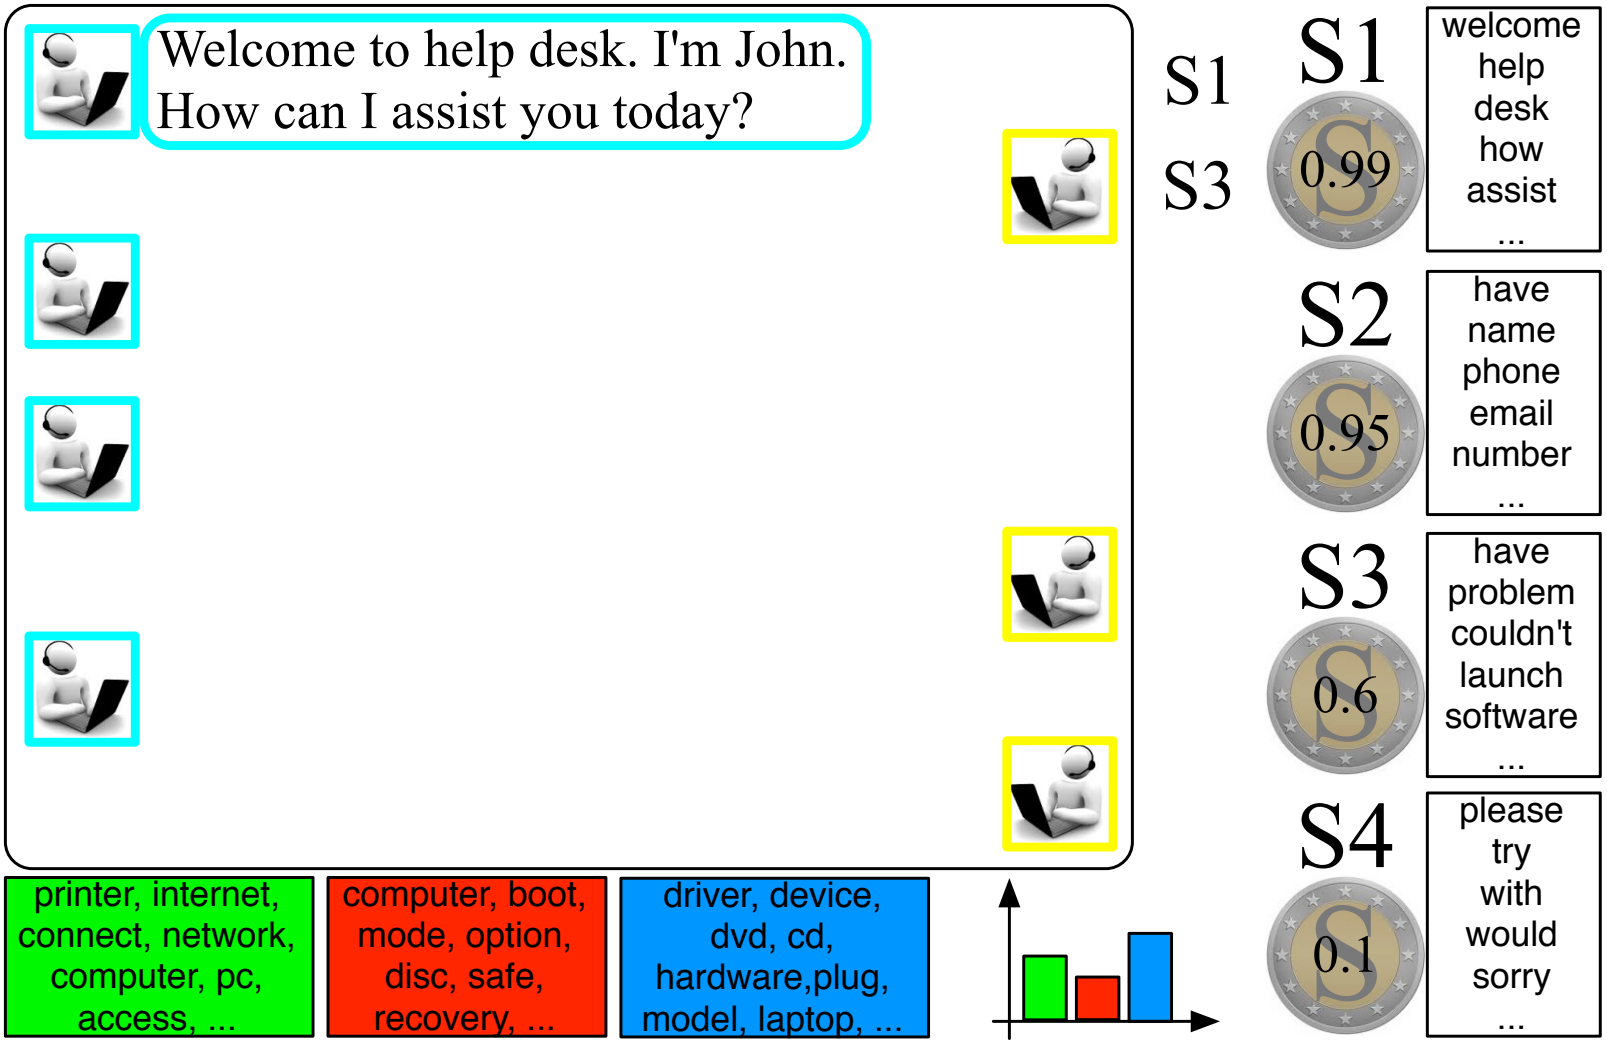

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

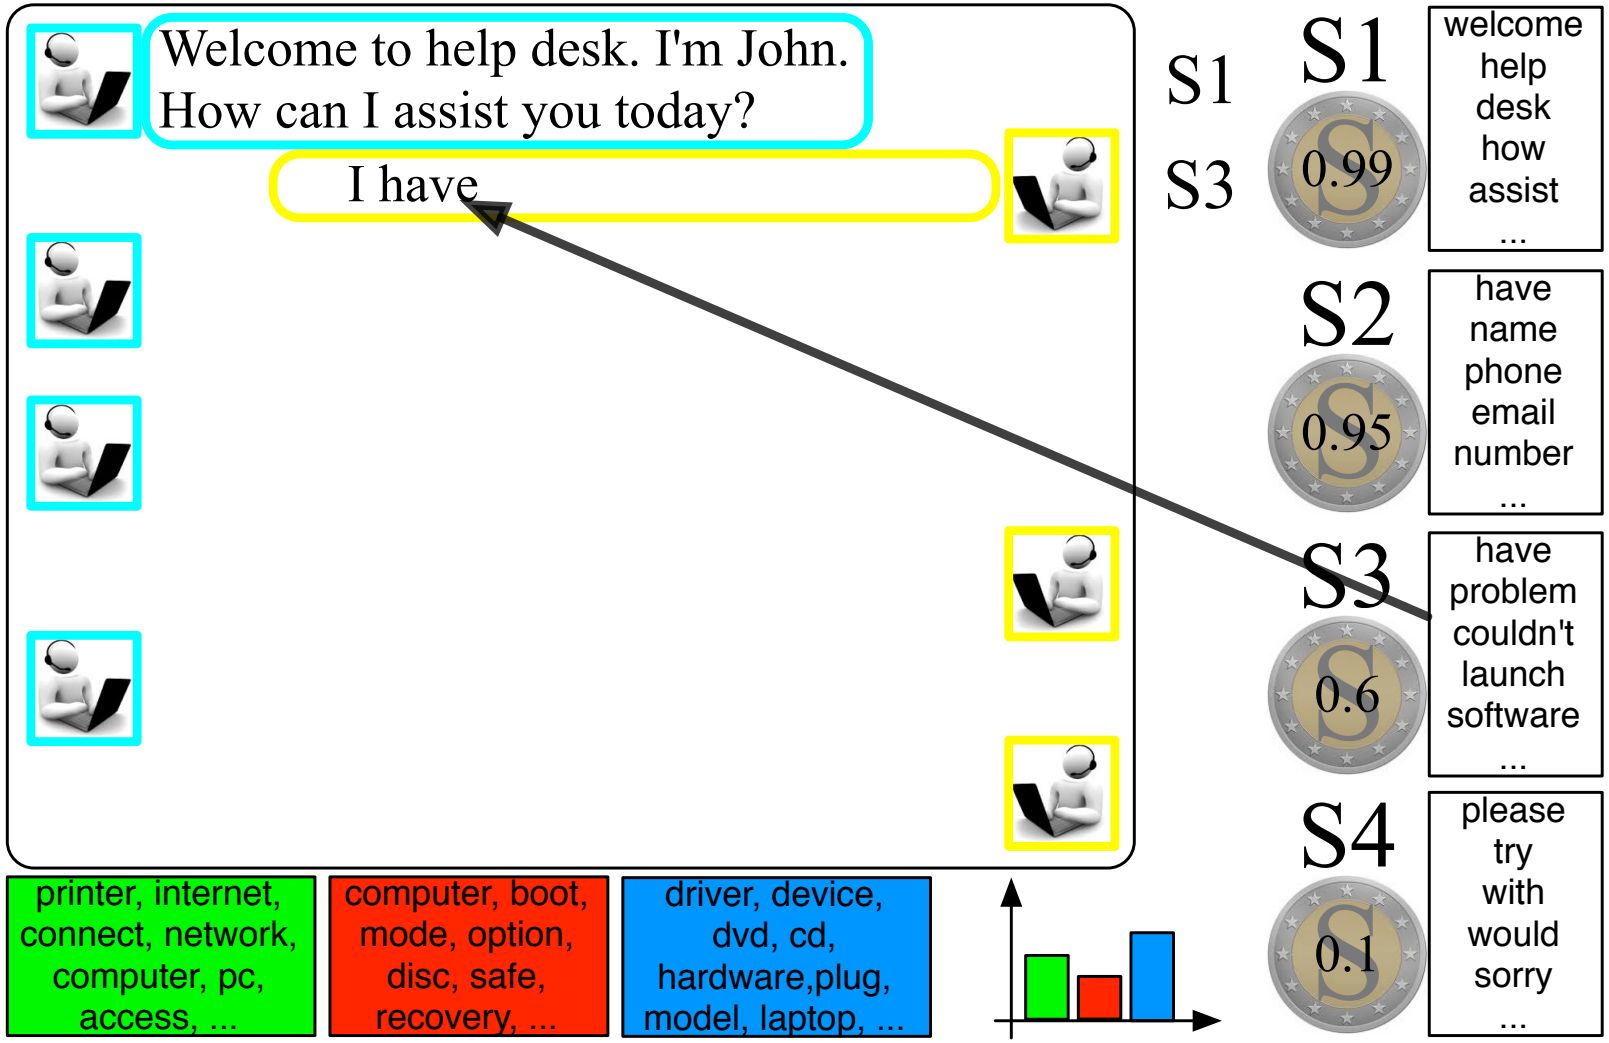

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

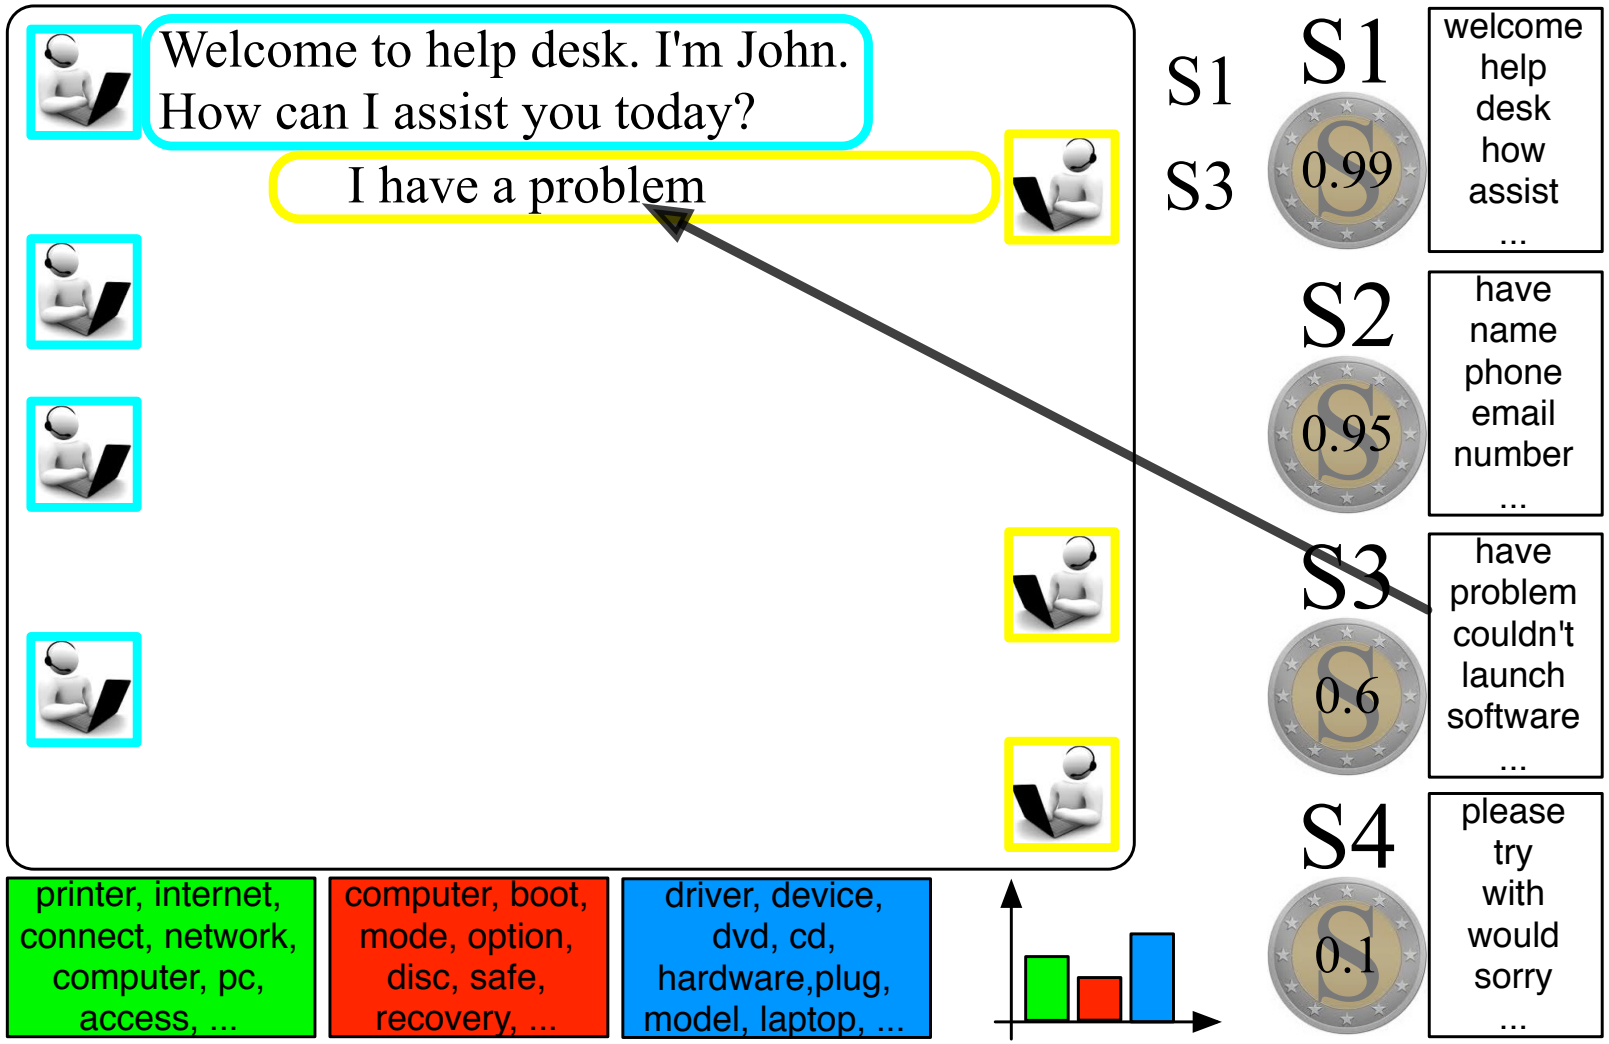

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

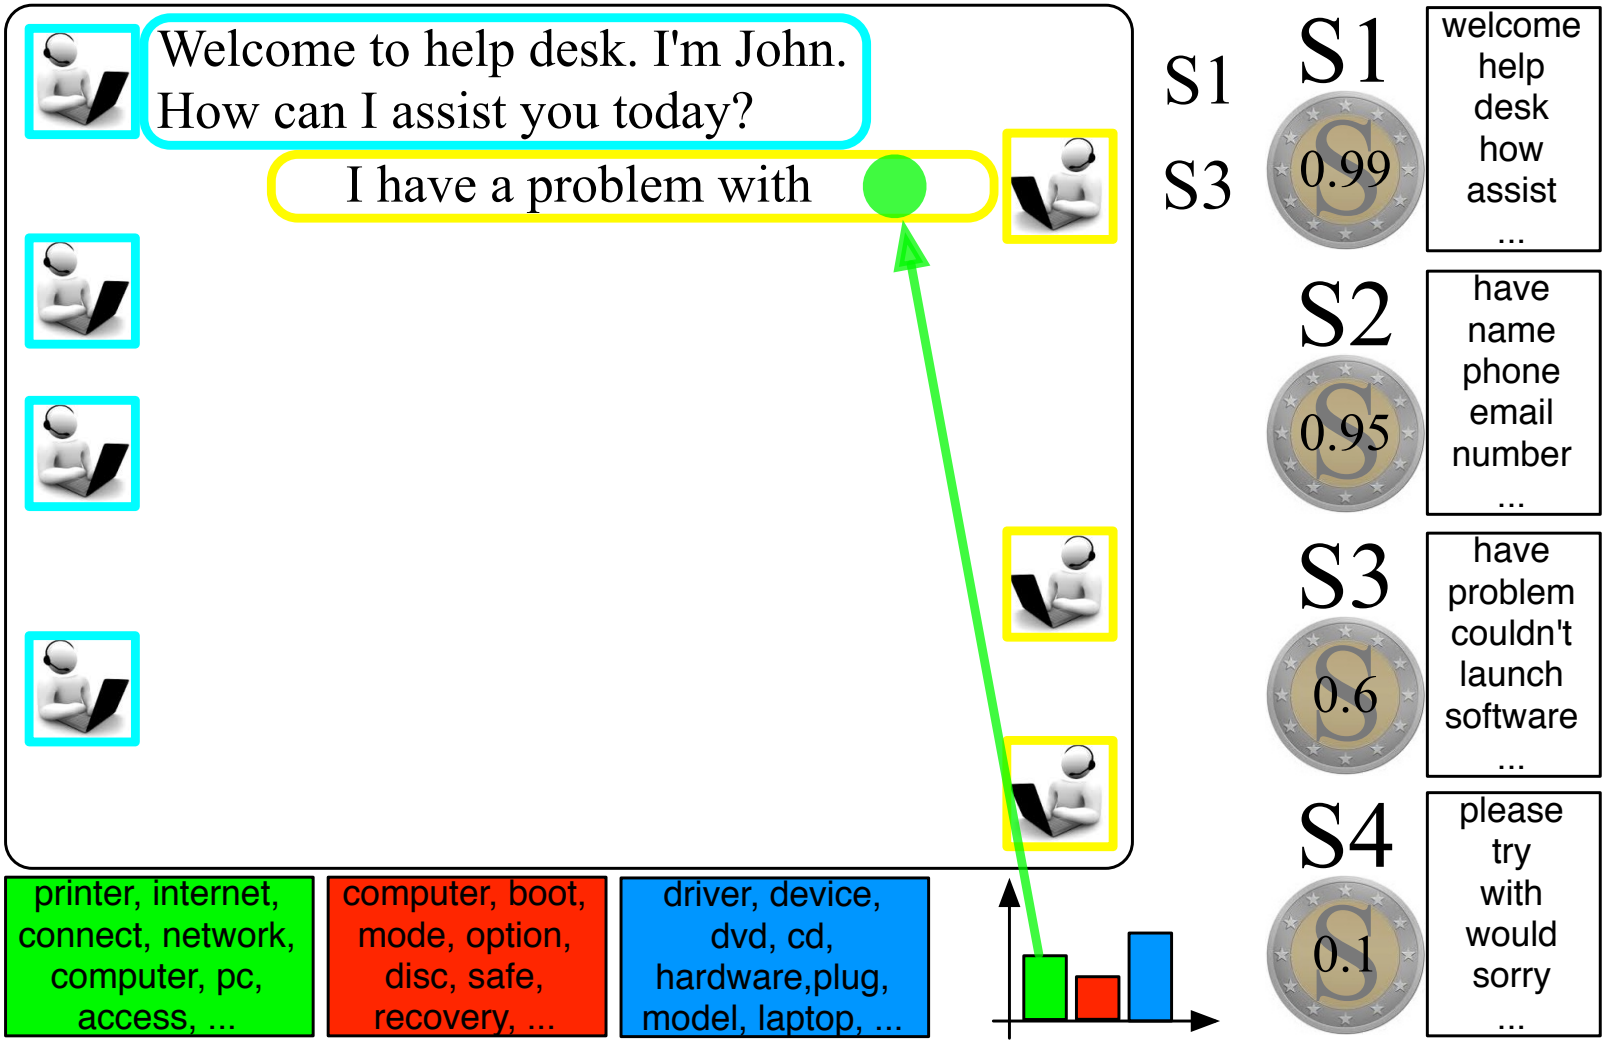

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

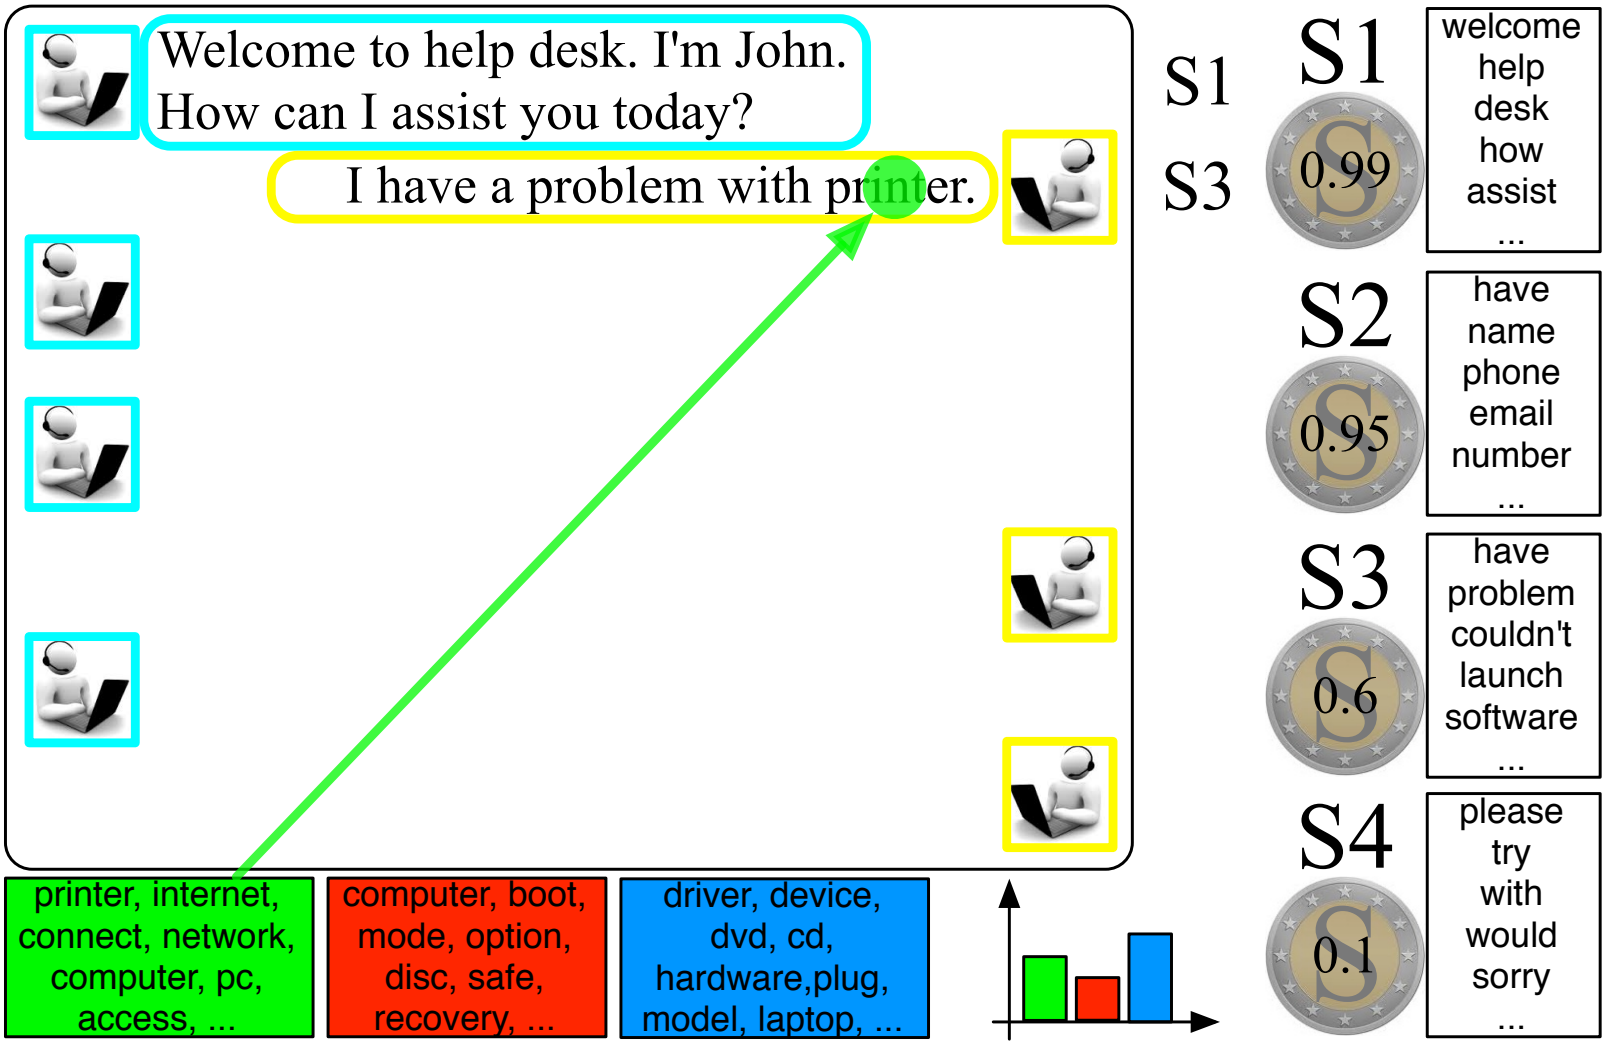

TM-HMMSS: Share Topics across Sessions

TM-HMMSS: HMM with state language model and state source generator and session topic model

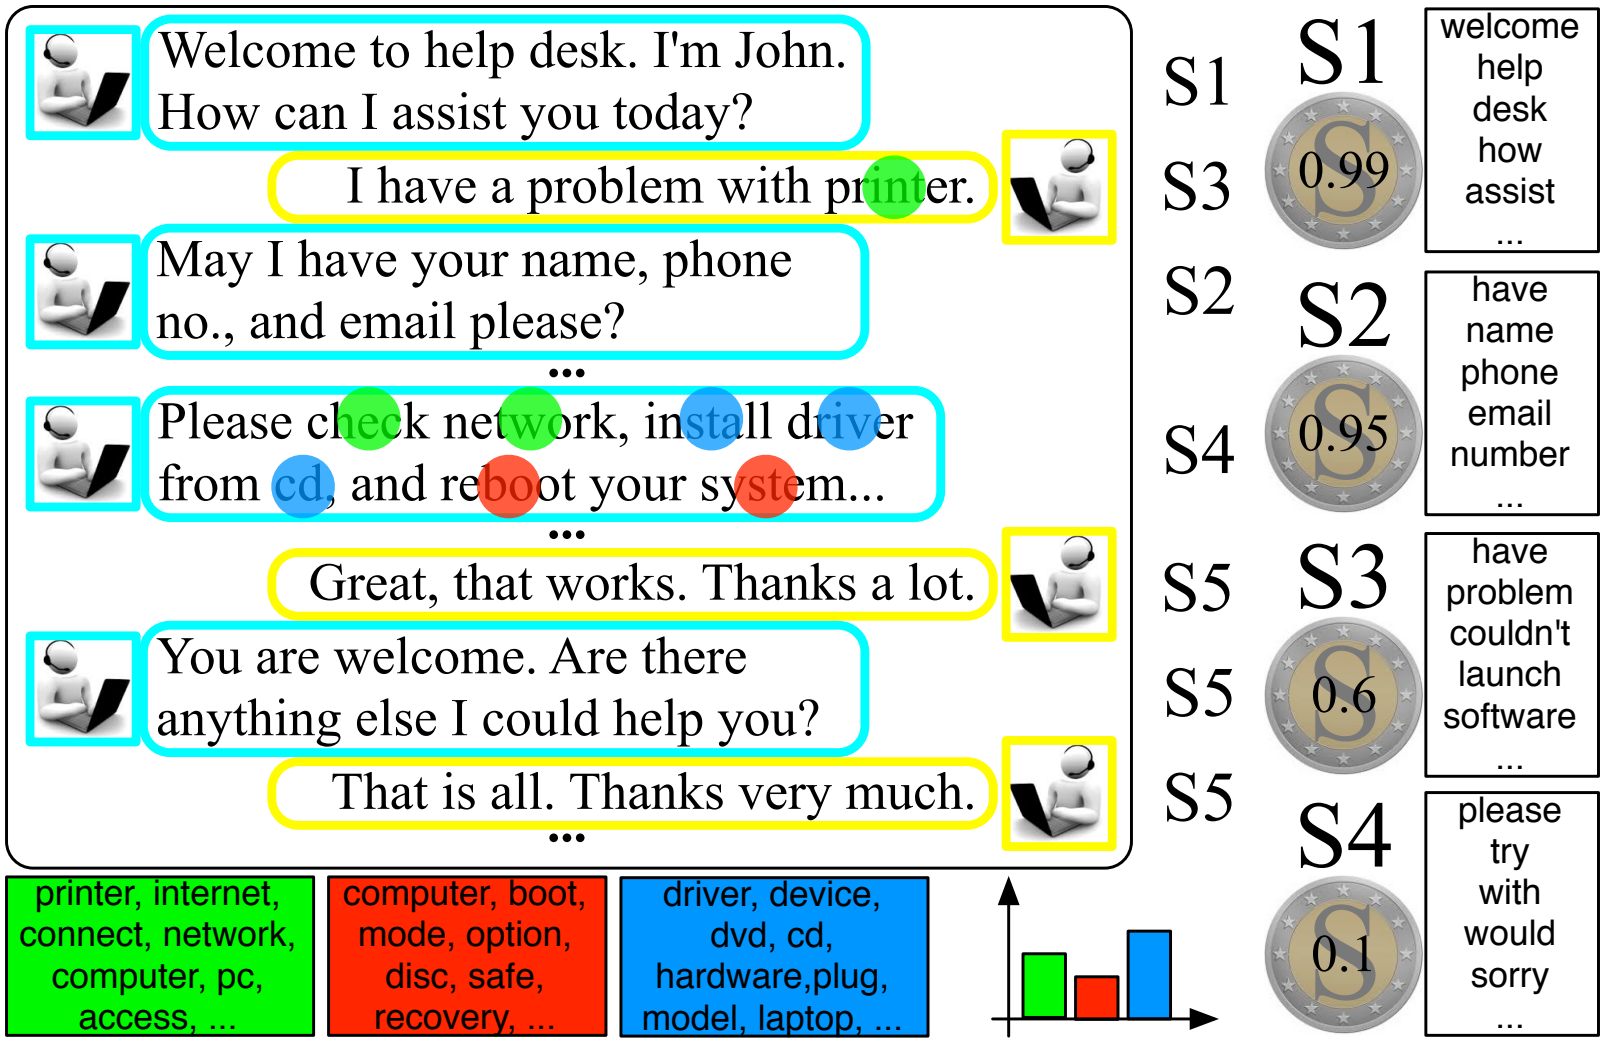

Outline

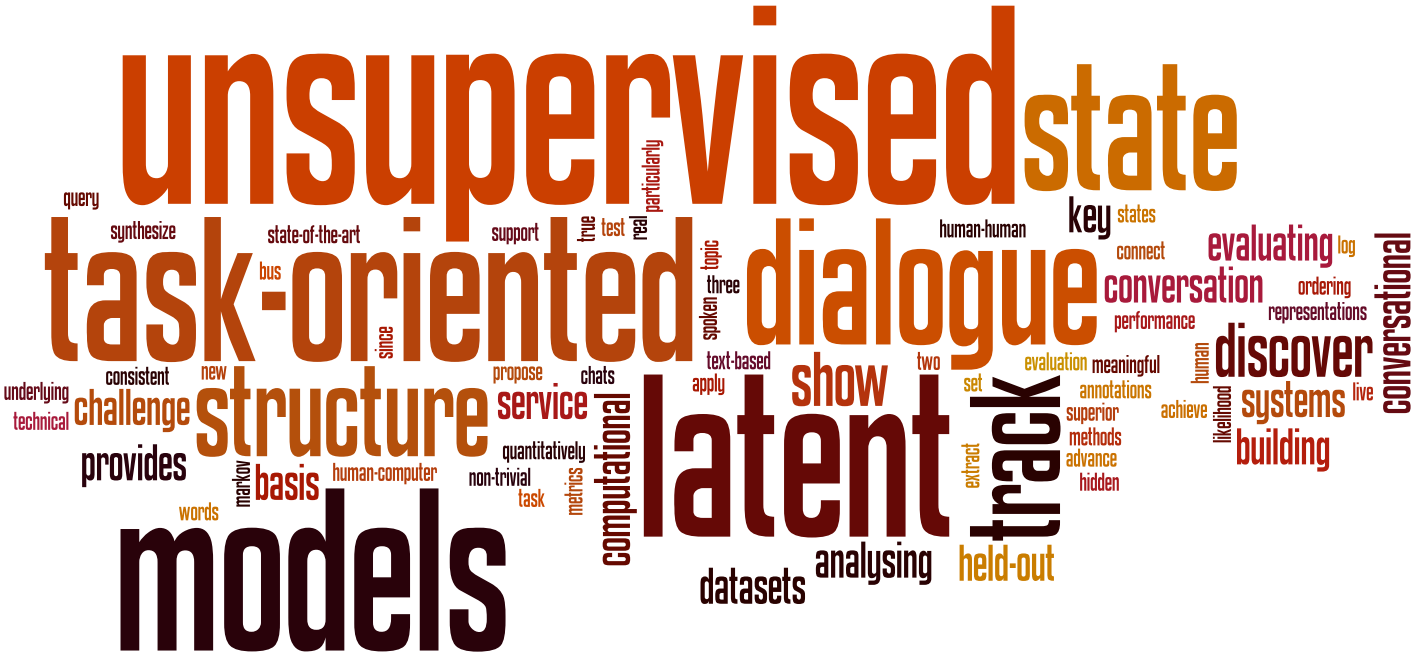

1. Existing Models: LM-HMM and LM-HMMS
2. Proposed Models: TM-HMM and TM-HMMSS
3. Experiments: qualitative and quantitative evaluations
4. Conclusion: what's exciting in Zhai & Williams (2014)...

Inferred Flowchart on *BusTime* Dataset

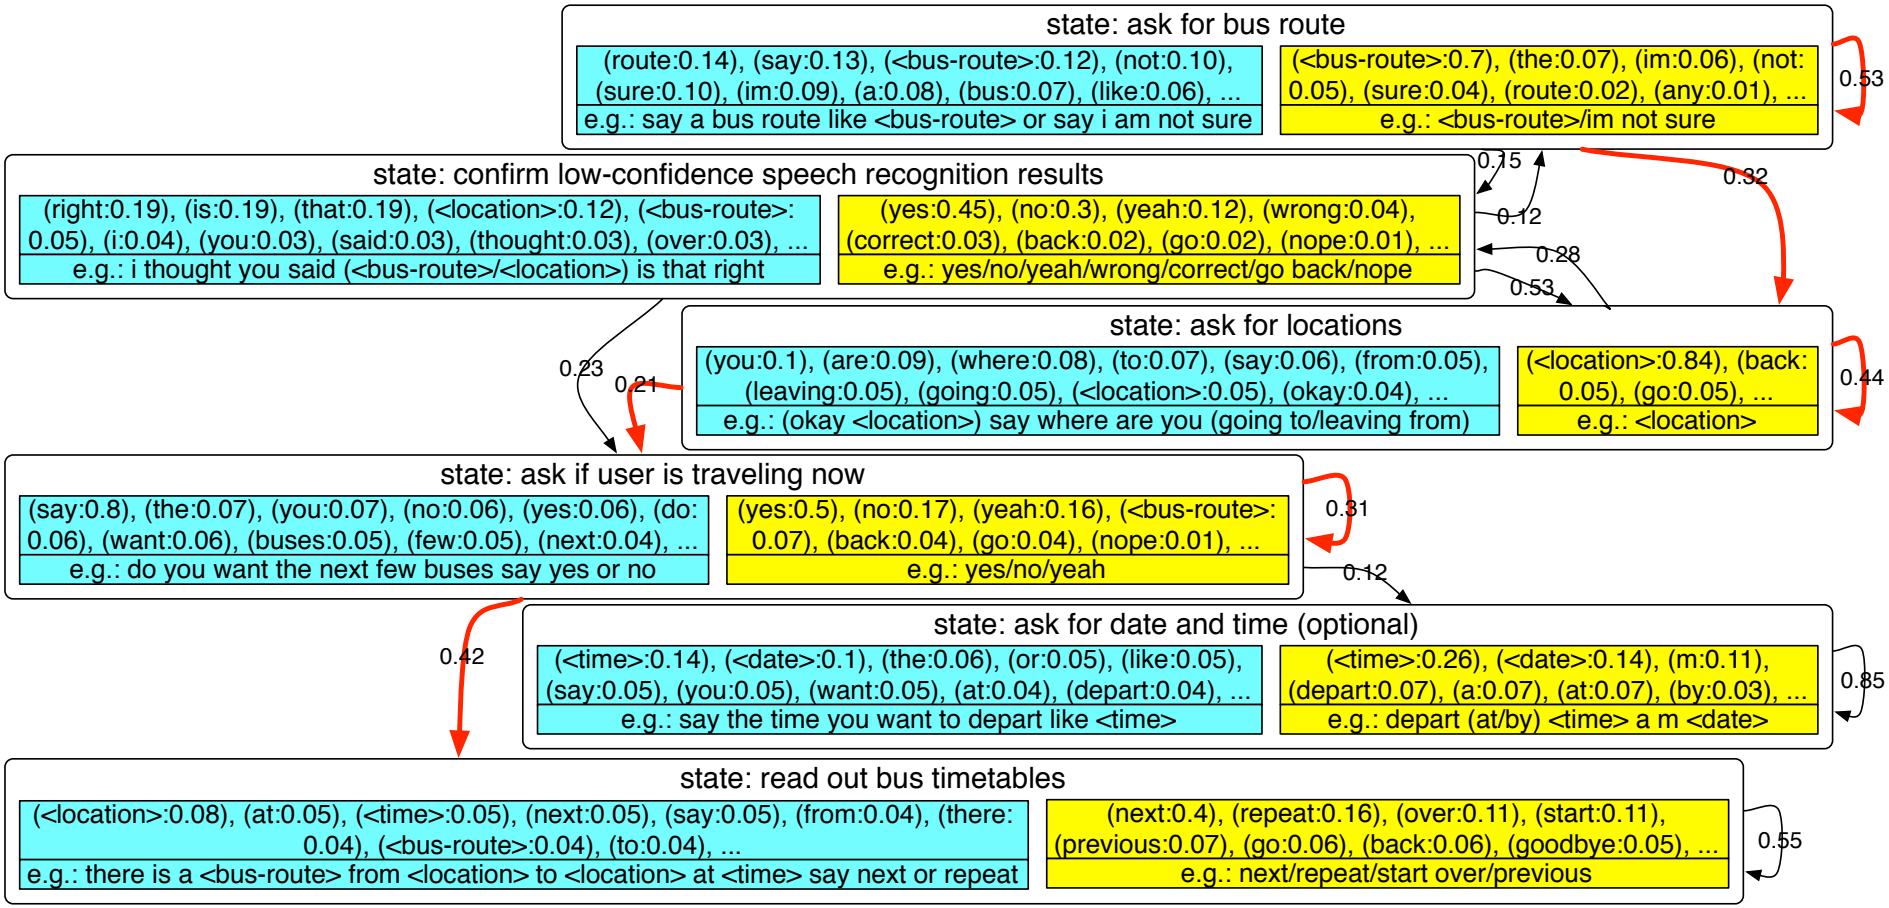

TM-HMM with 10 states and 10 topics
(transition cutoff threshold is 0.1)

Inferred Flowchart on *BusTime* Dataset

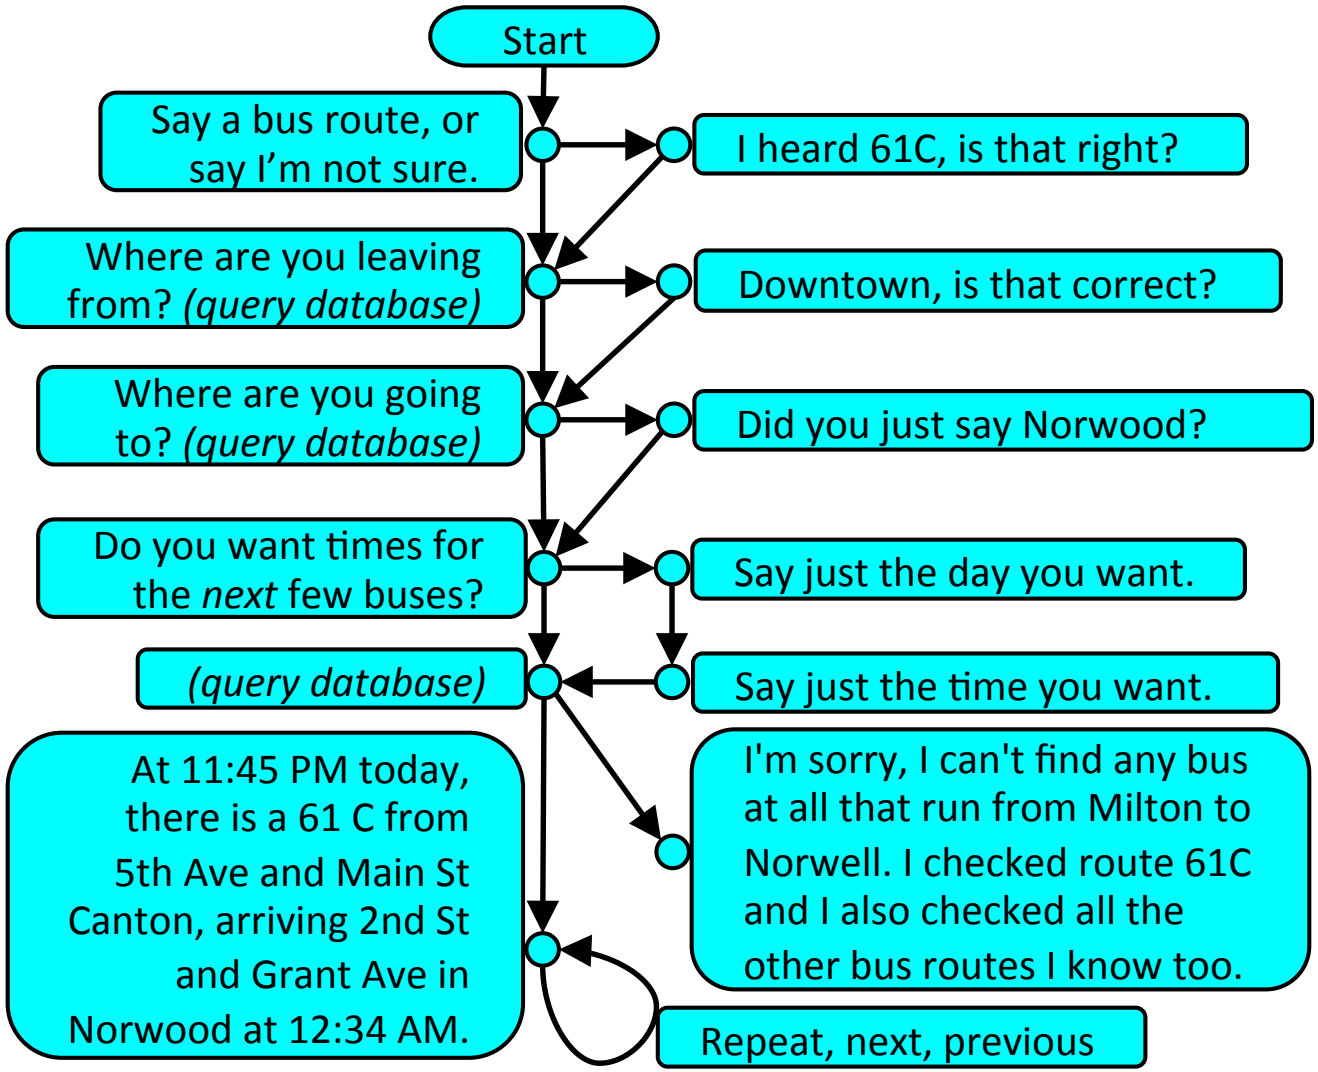

Hand-crafted reference flowchart for *BusTime* (Williams, 2012)

Inferred Flowchart on *TechSupport* Dataset

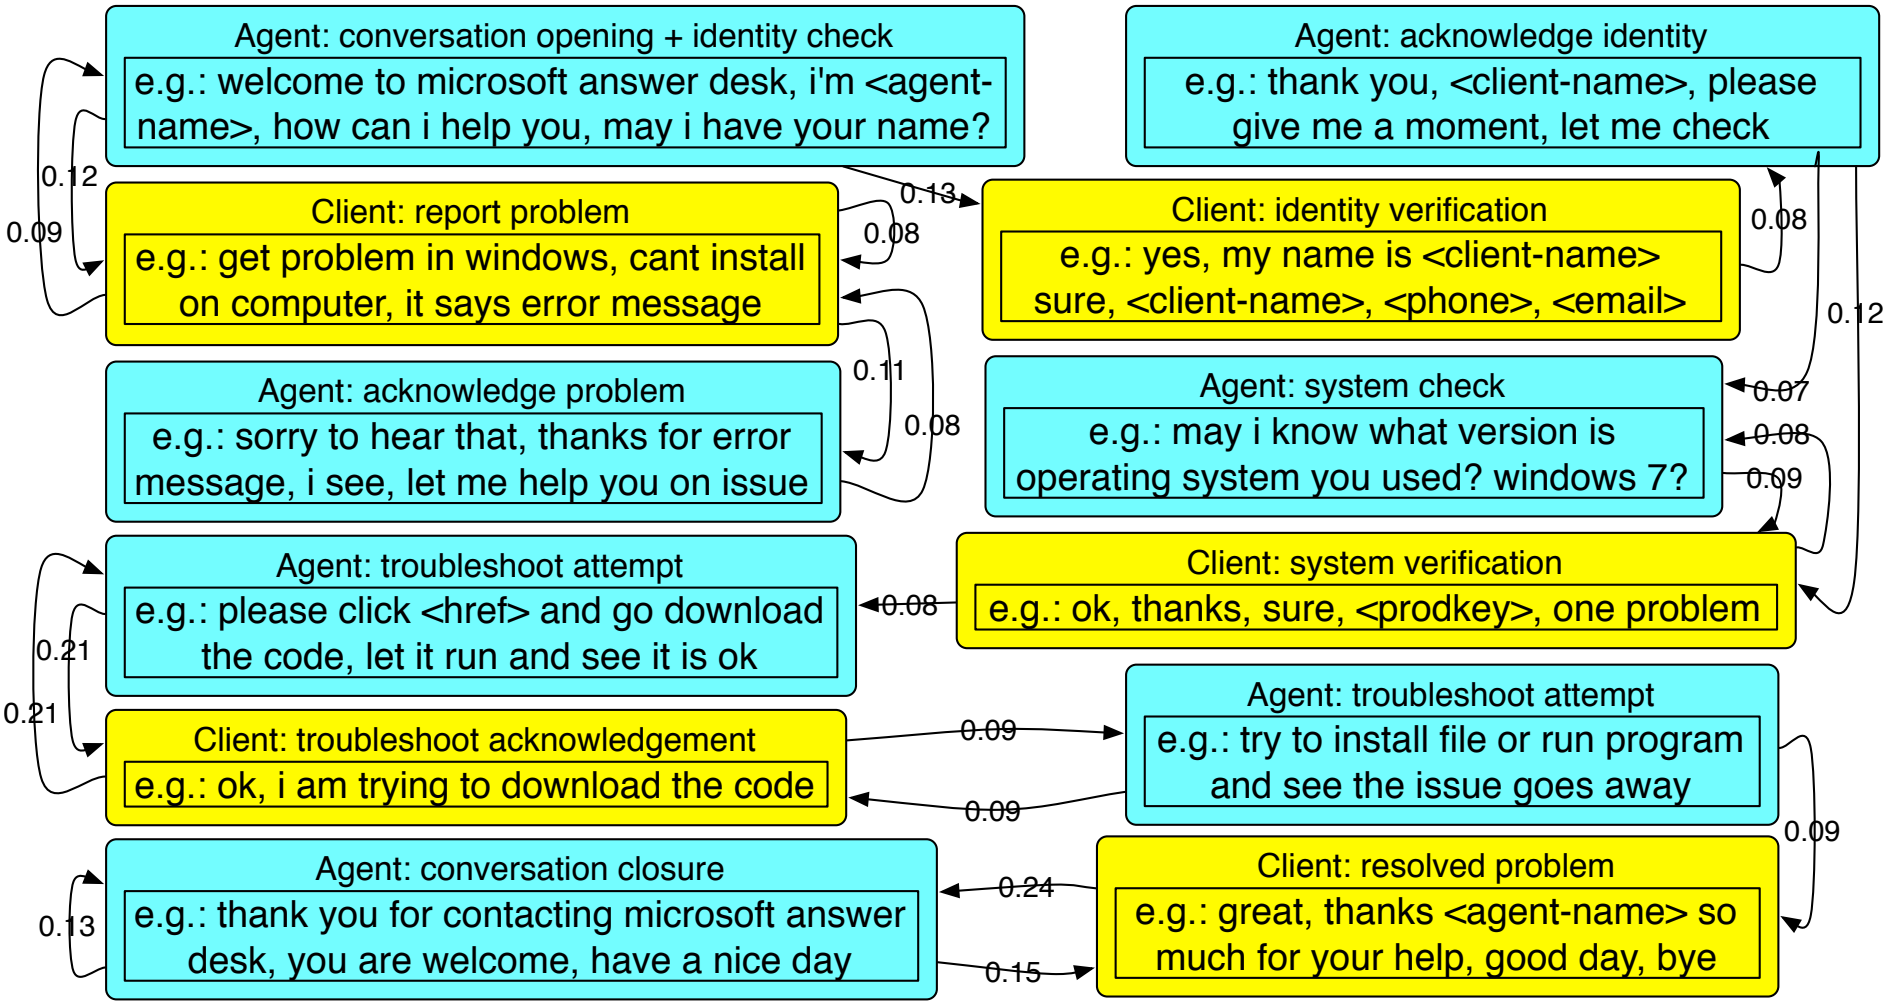

TM-HMMSS with 20 states and 20 topics
(transition cutoff threshold is 0.05)

Inferred Topics

Topic	Top Ranked Words
browser	internet, explor, browser, ie, open, websit, googl, download, click, chrome, . . .
backup	file, restor, system, comput, back, folder, creat, option, dont, delet, . . .
boot	comput, boot, mode, option, disc, safe, recoveri, repair, back, clean, cd, disk, . . .
update	updat, window, servic, instal, pack, run, comput, download, check, restart, inform, system, error, fix, . . .
network	connect, internet, printer, comput, network, pc, print, access, wireless, hp, cable, adapt, router, speed, . . .
anti-virus	viru, scan, comput, remov, secur, run, system, anti, essenti, infect, defend, softwar, program, protect, antiviru, malwar, . . .
hardware	driver, devic, drive, dvd, cd, hardwar, issu, model, laptop, plug, software, usb, . . .
windows	window, upgrad, 8, download, 7, instal, bit, vista, pro, system, . . .
office	offic, 2010, word, microsoft, home, excel, version, 2007, student, document, trial, 2013, . . .

TM-HMMSS with 20 states and 20 topics
(transition cutoff threshold is 0.05)

Log-likelihood Evaluation on *BusTime* Dataset

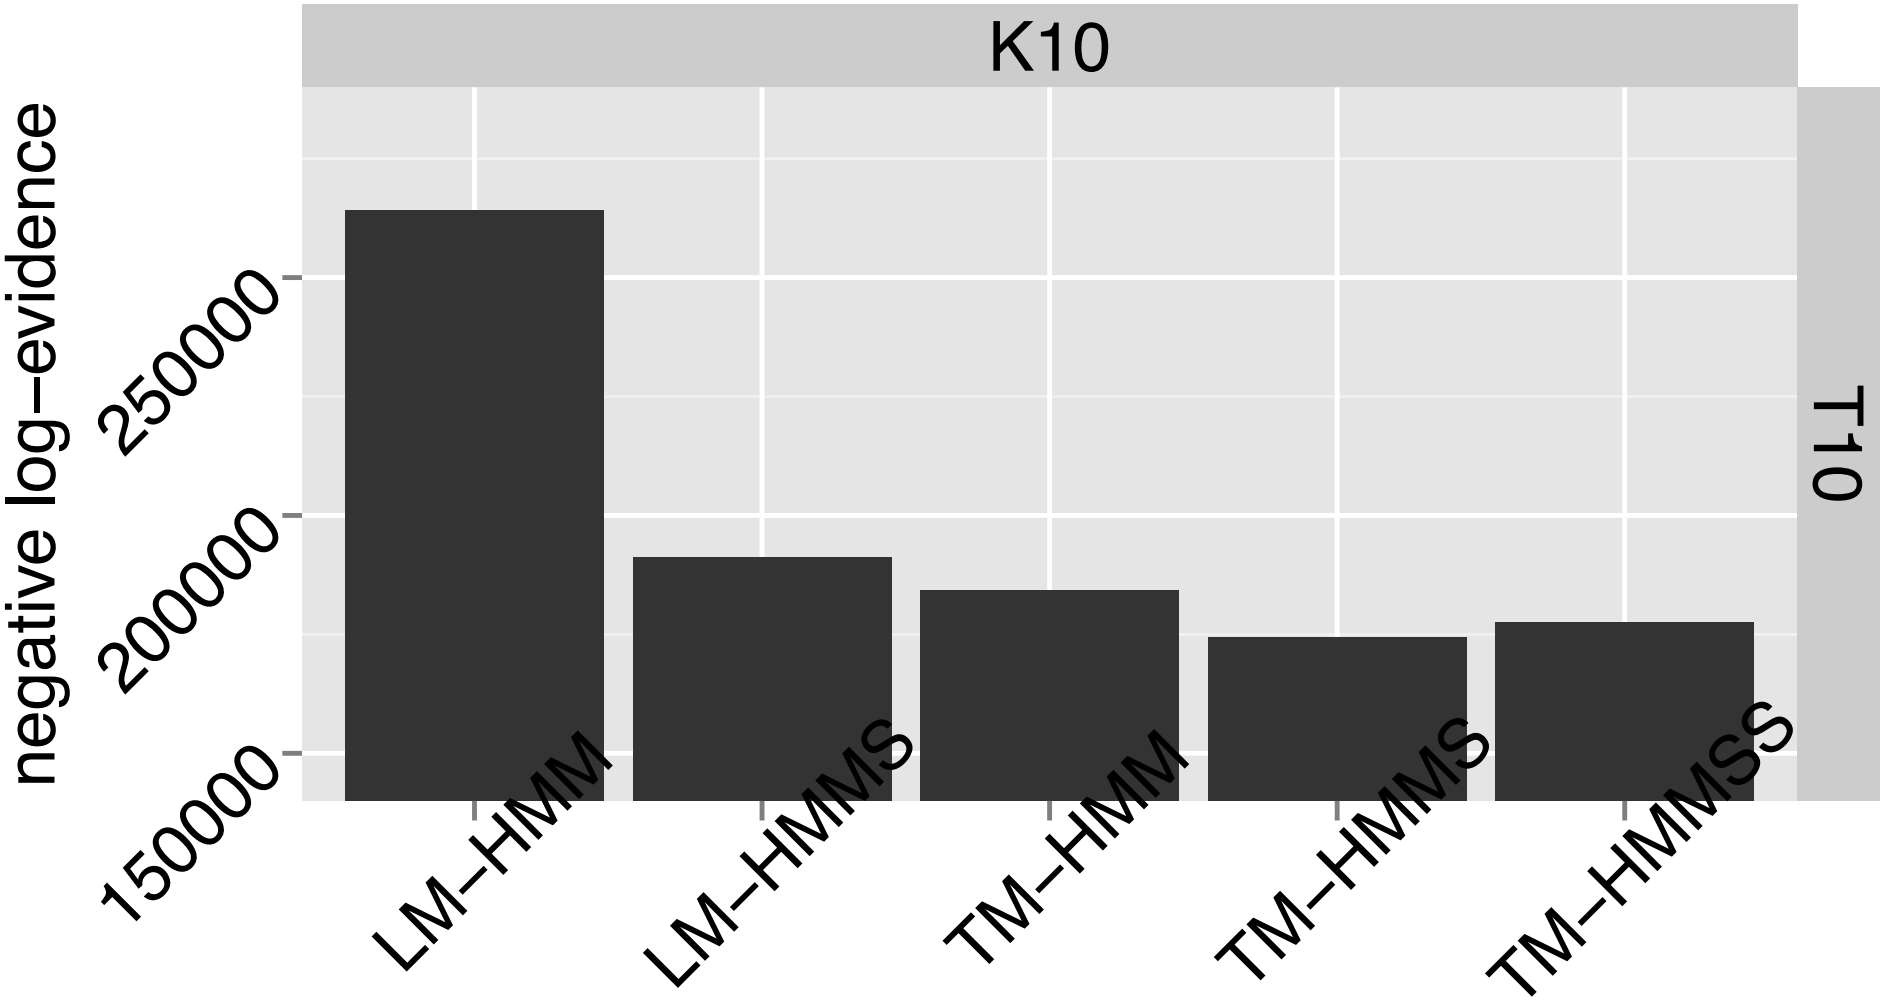

Negative log-likelihood (lower the better) on held-out test set with 10 states and 10 topics

Log-likelihood Evaluation on *TechSupport* Dataset

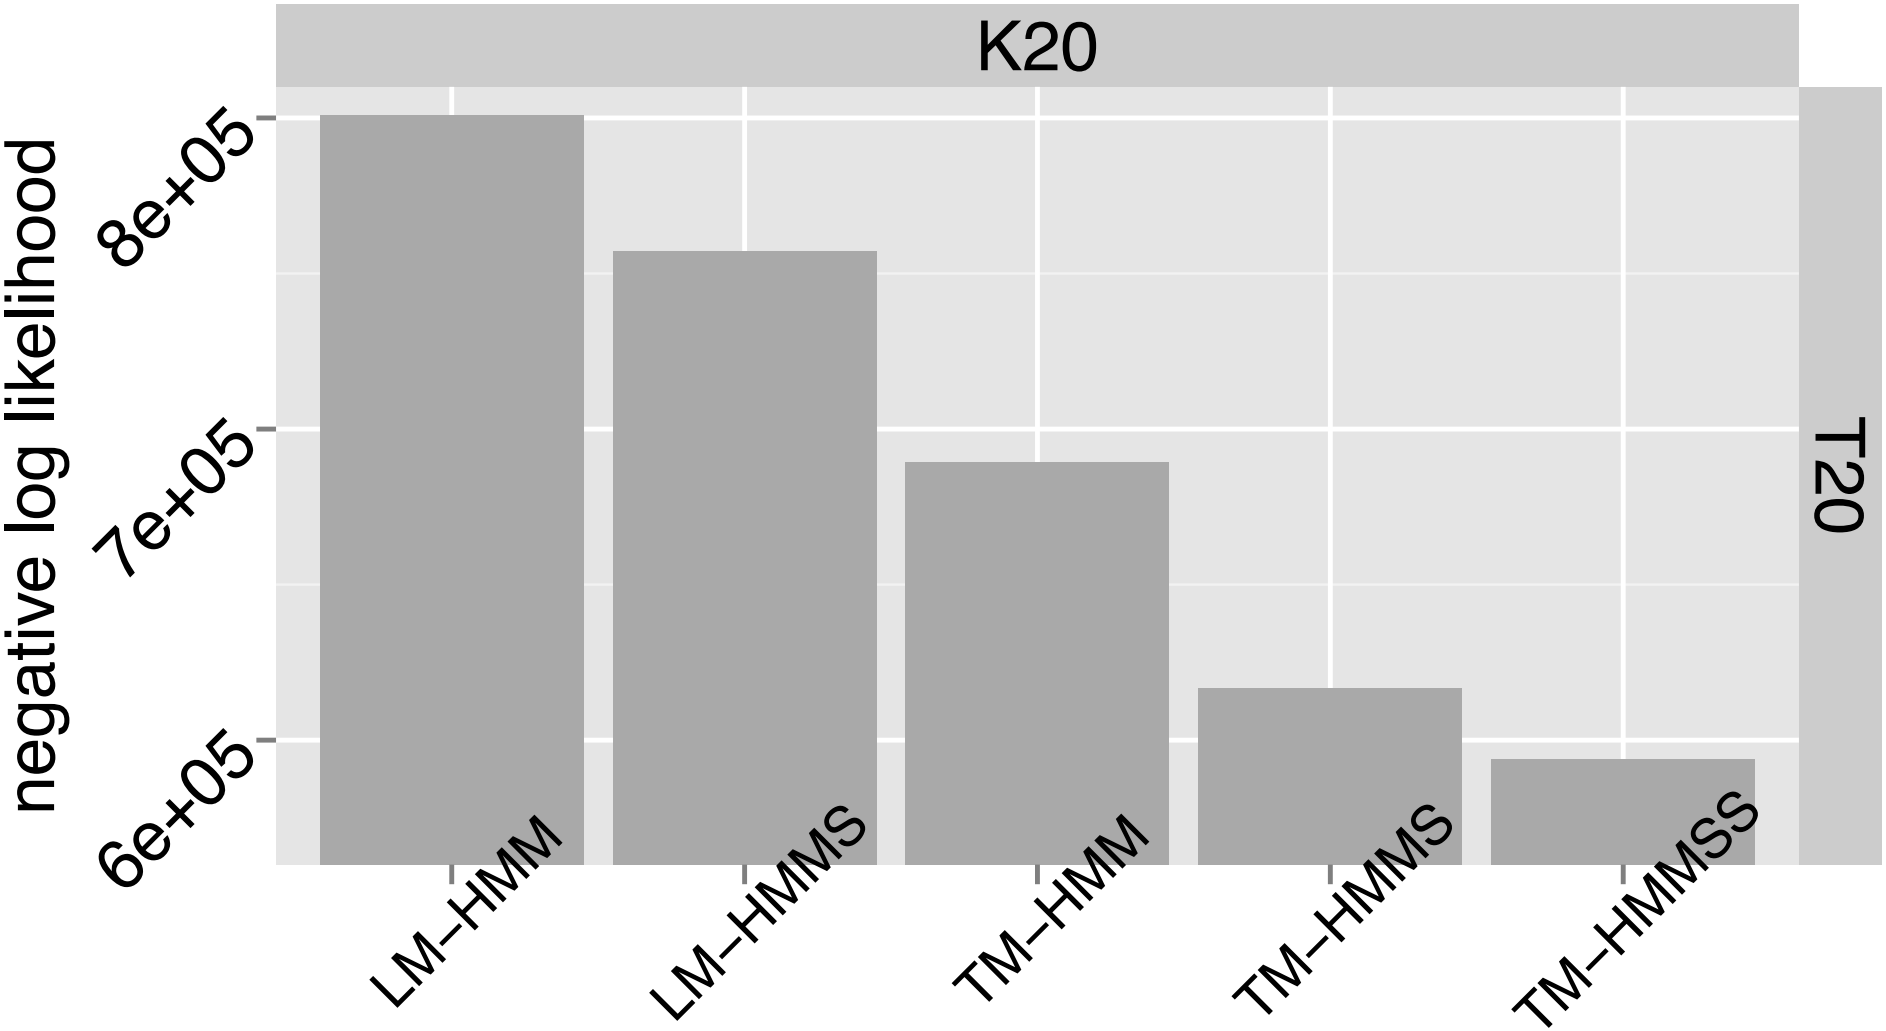

Negative log-likelihood (lower the better) on held-out test set with 20 states and 20 topics

Ordering Task

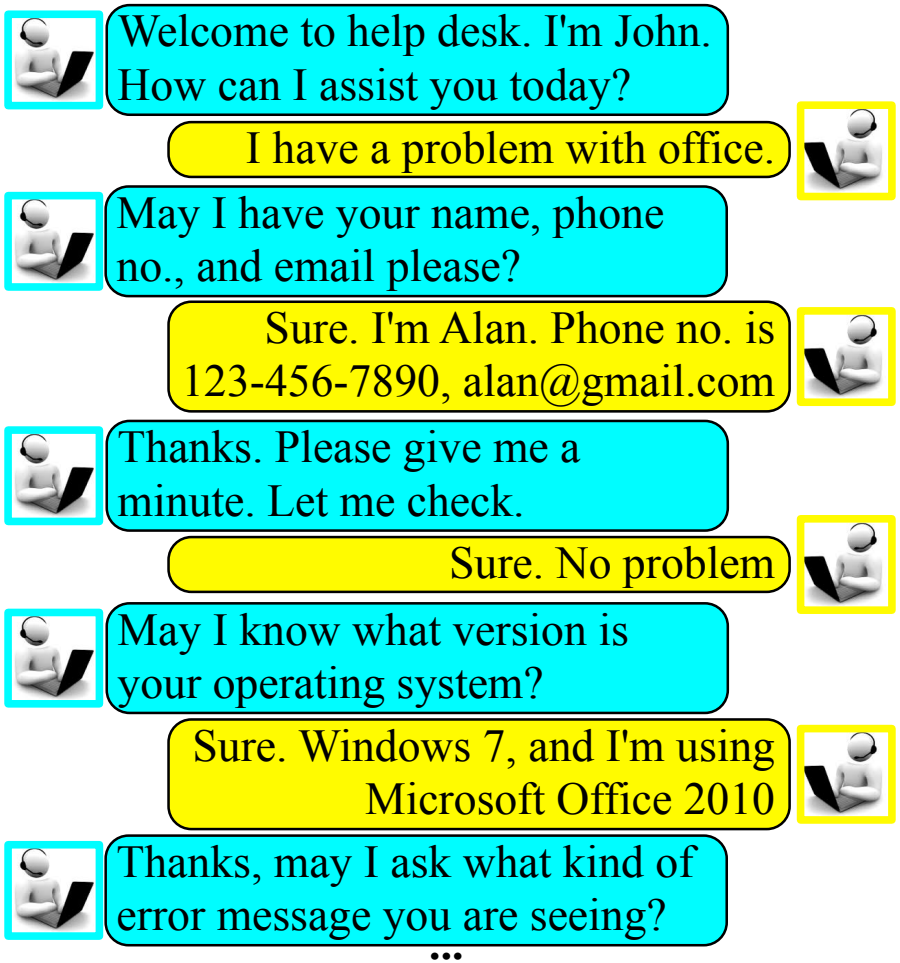

Ordering Task

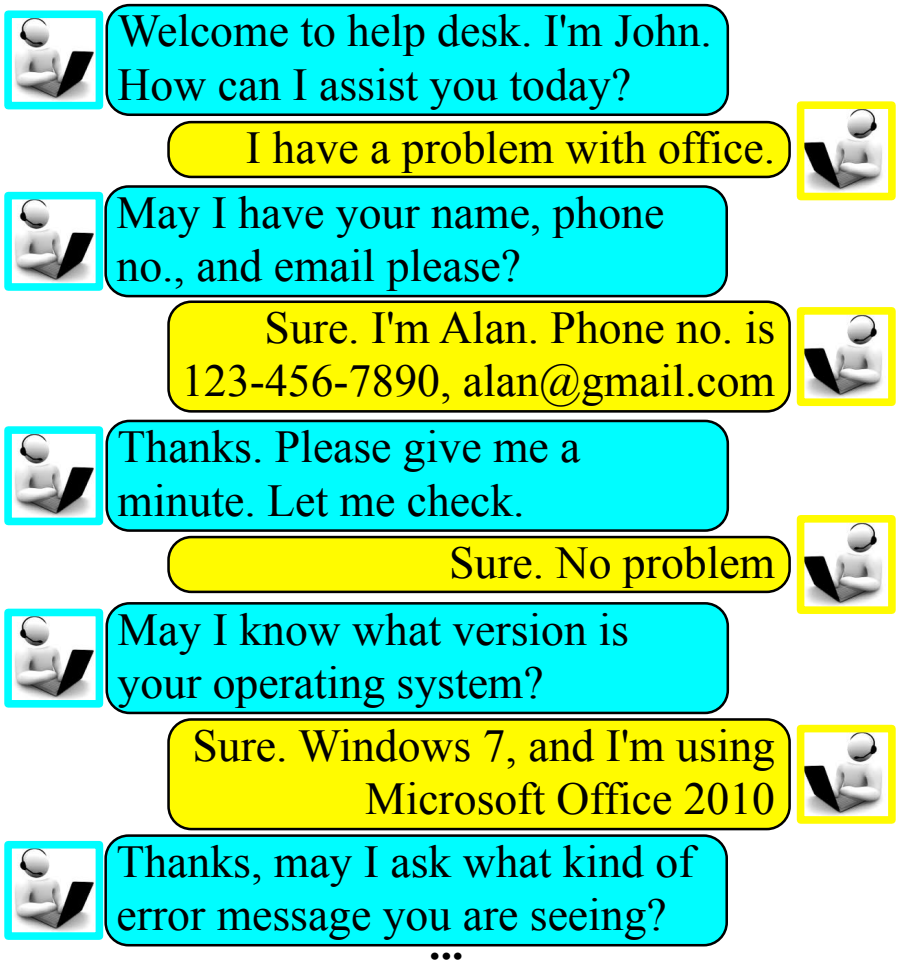

Ordering Task

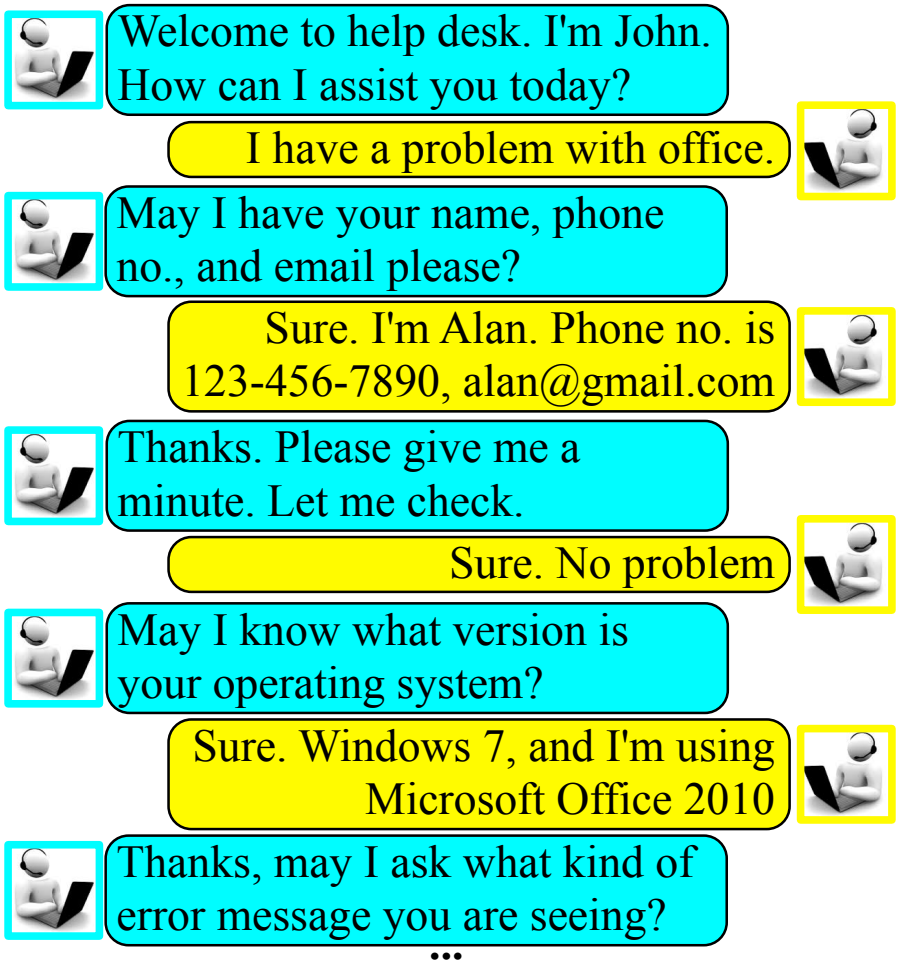

$-1 \leq$ correlation $\leq +1$

reverse ordering — identical ordering

Ordering Task on *BusTime* Dataset

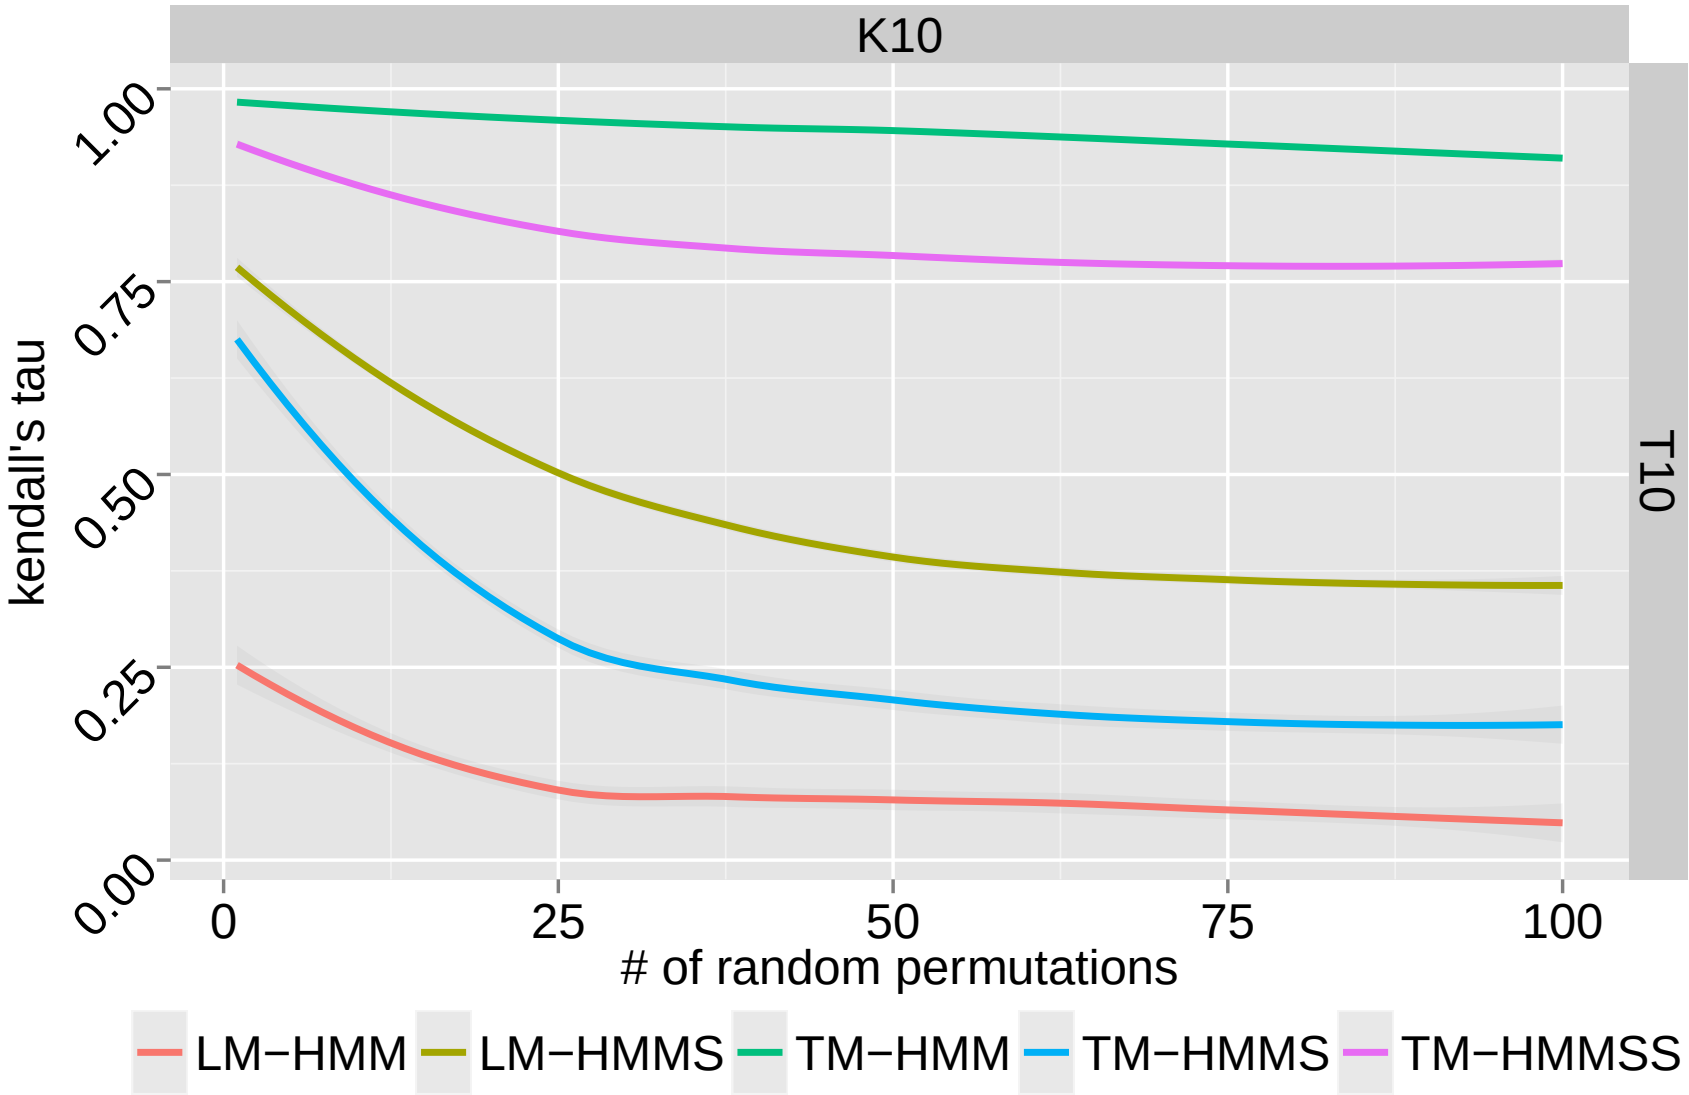

Ordering Task on *TechSupport* Dataset

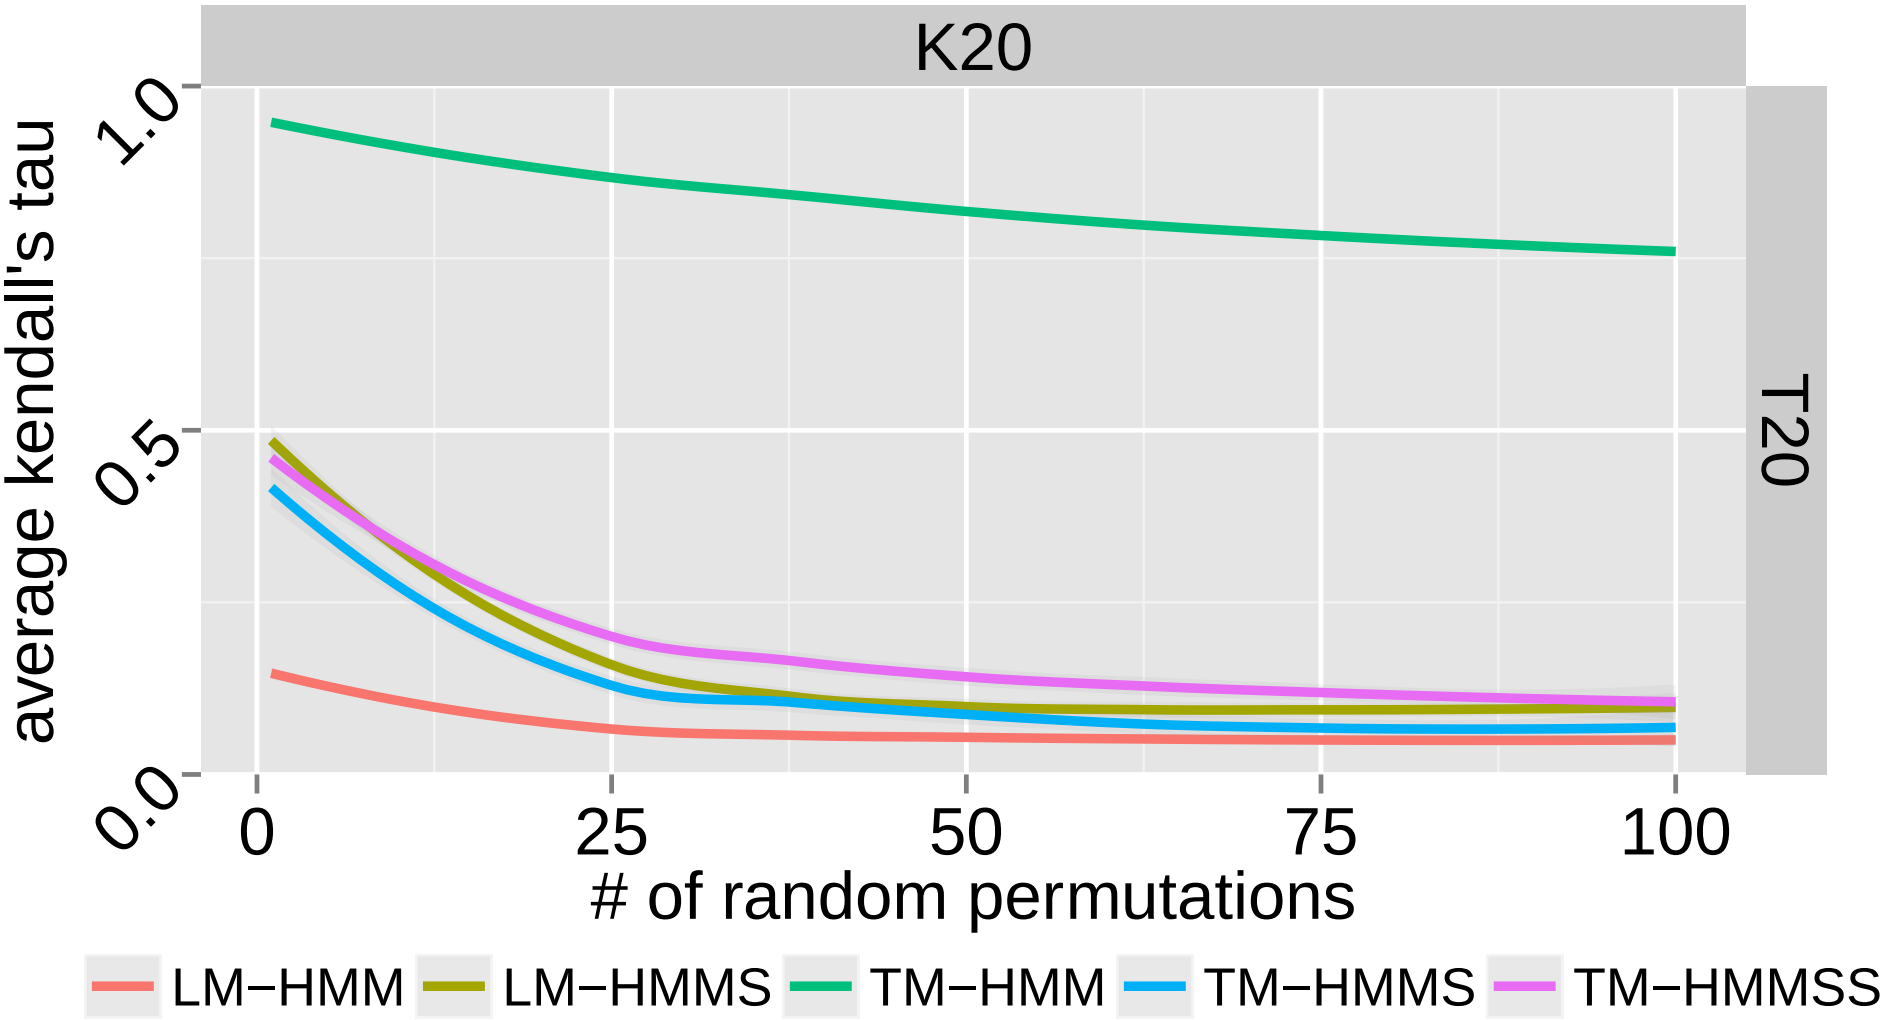

Outline

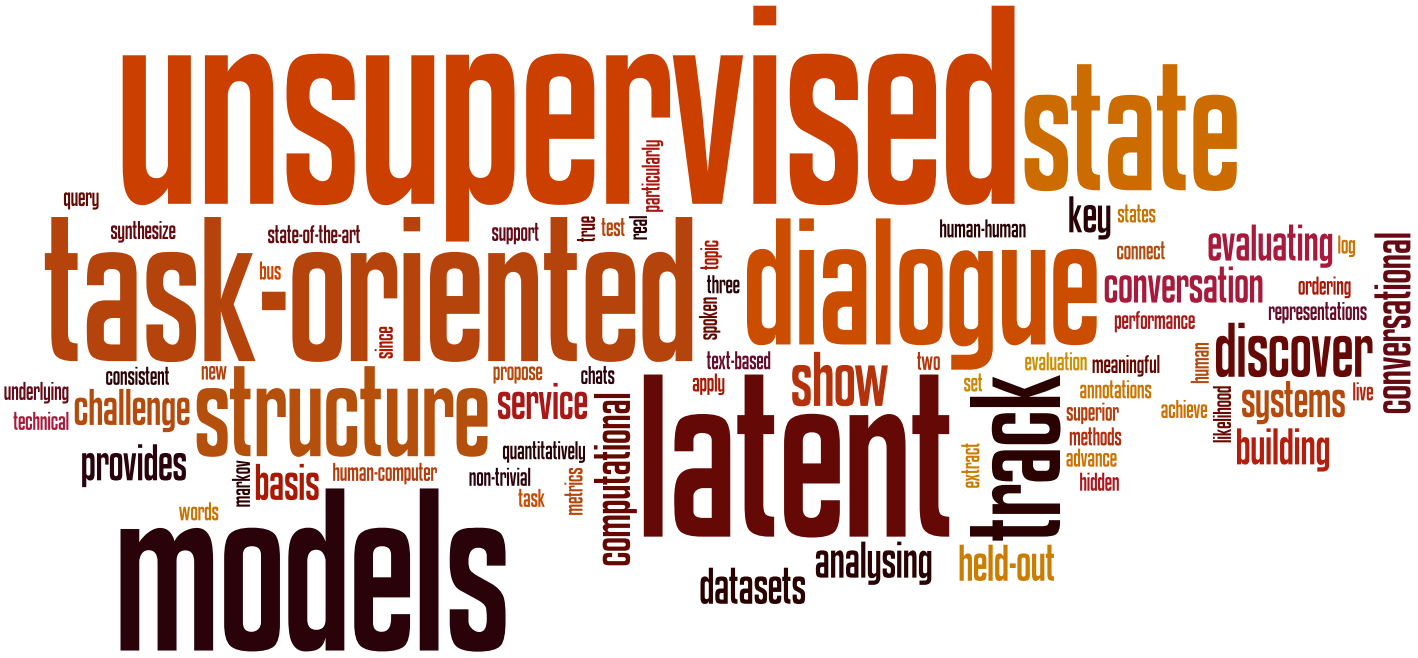

1. Existing Models: LM-HMM and LM-HMMS
2. Proposed Models: TM-HMM and TM-HMMSS
3. Experiments: qualitative and quantitative evaluations
4. Conclusion: what's exciting in Zhai & Williams (2014)...

Conclusion

In this talk

1. synthesize hidden Markov models and topic models
2. propose unsupervised generative models: TM-HMM and TM-HMMSS
3. discover two horizons of a dialogue system automatically
   1) a generic flow-chart and 2) a domain of topics

What's exciting in Zhai & Williams (2014)...

- Plate notations, generative stories and MCMC inference
- Experiments on *BusTime* and *TechSupport* datasets
  - *BusTime*: transcripts from human-computer interactive system
  - *TechSupport*: a set of human-computer online chat dialogues
  - Qualitative: meaningful dialogue flows, aligned with reference design
  - Quantitative: log-likelihood and ordering task on held-out test set
- Sensitivity to different parameter settings

The End!

Thank you!

Questions?

Special Acknowledgement
Alan Ritter, Bill Dolan
Jordan Boyd-Graber
Kai (Anthony) Lui

Bibliography I

Barzilay, Regina and Lee, Lillian. Catching the drift: Probabilistic content models, with applications to generation and summarization. In *Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 113–120, 2004.

Ritter, Alan, Cherry, Colin, and Dolan, Bill. Unsupervised modeling of twitter conversations. In *Conference of the North American Chapter of the Association for Computational Linguistics*, 2010.

Williams, Jason D. Challenges and opportunities for state tracking in statistical spoken dialog systems: Results from two public deployments. *Journal of Selected Topics in Signal Processing*, 2012.

Zhai, Ke and Williams, Jason D. Discovering latent structure in task-oriented dialogues. In *Proceedings of the Association for Computational Linguistics*, 2014.