

Hey, Lumi!

Using Natural Language for Intent-based Network Management

Arthur S. Jacobs¹

Ricardo J. Pfitscher¹

Rafael H. Ribeiro¹

Ronaldo A. Ferreira²

Lisandro Z. Granville¹

Walter Willinger³

Sanjay G. Rao⁴

July 15th, 2021

Deploying network policies is hard...

Deploying network policies is hard...

“Block YouTube in the office”

Deploying network policies is hard...

“Block YouTube in the office”

“Inspect all traffic for student dorms.”

Deploying network policies is hard...

“Block YouTube in the office”

“Inspect all traffic for student dorms.”

“Rate limit employees streaming traffic”

Deploying network policies is hard...


“Block YouTube in the office”

“Inspect all traffic for student dorms.”

“Rate limit employees streaming traffic”




How to deploy network policies?




How to deploy network policies?

Higher-level



How to deploy network policies?

Higher-level





What if we use natural language?

Lumi 

Lumi

1. Allows network operators to express their high-level intents in **Natural Language**

Lumi

1. Allows network operators to express their high-level intents in **Natural Language**
2. Relies on a **high-level intent language** to ask for confirmation on intents before deployment

Lumi

1. Allows network operators to express their high-level intents in **Natural Language**
2. Relies on a **high-level intent language** to ask for confirmation on intents before deployment
3. **Learns over time** using operator knowledge

Lumi

1. Allows network operators to express their high-level intents in **Natural Language**
2. Relies on a **high-level intent language** to ask for confirmation on intents before deployment
3. **Learns over time** using operator knowledge
4. Evaluated using **real-world intents** from US University network policies, and a **User Study**

Lumi

Supported policies

ACL

QoS

Middlebox chaining

Temporal behavior

Related work

Network Configuration

Learns through Feedback



Hey network, can you understand me?

Lumi 

—



—

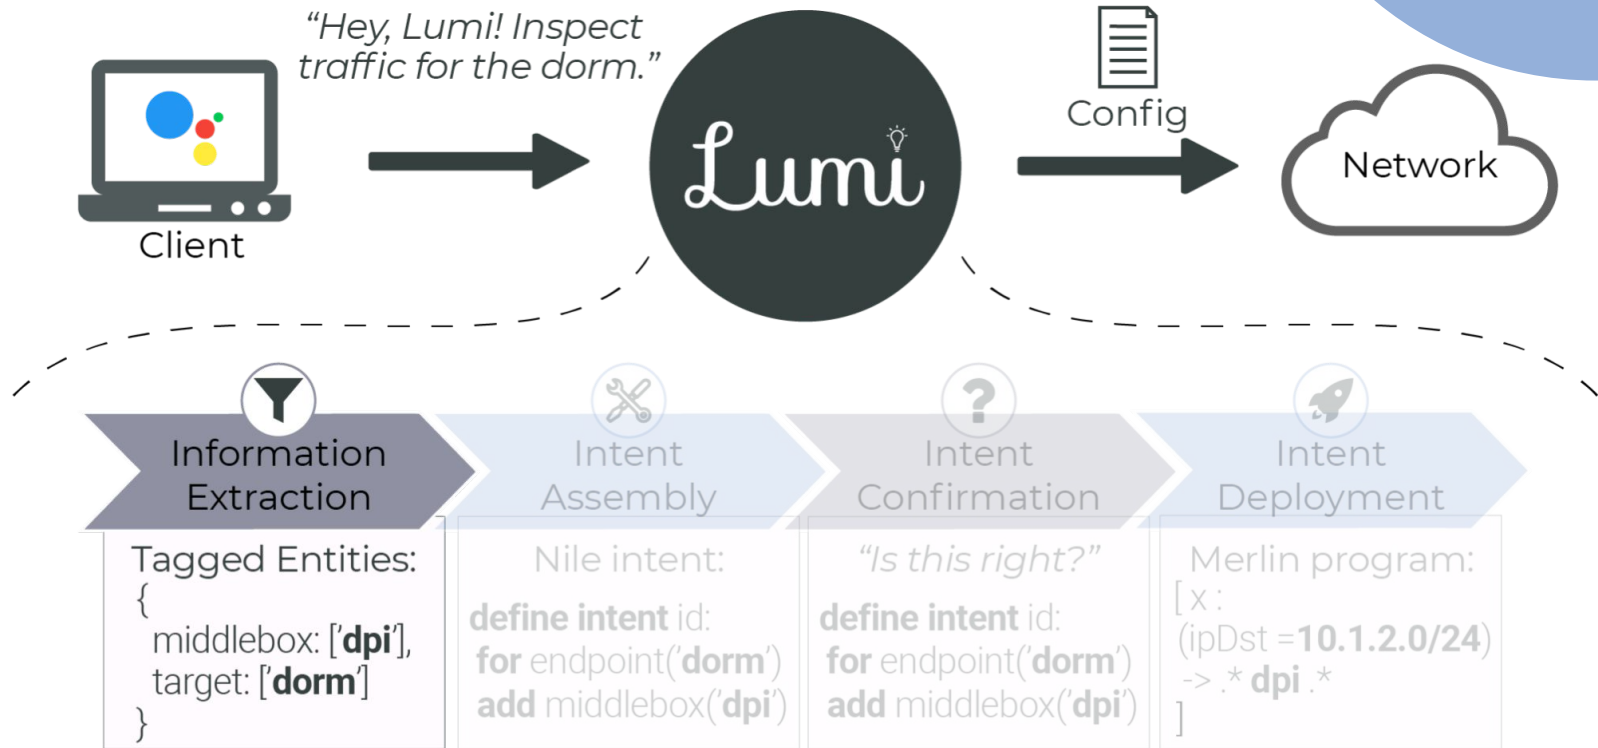
—



Lumi in a Nutshell

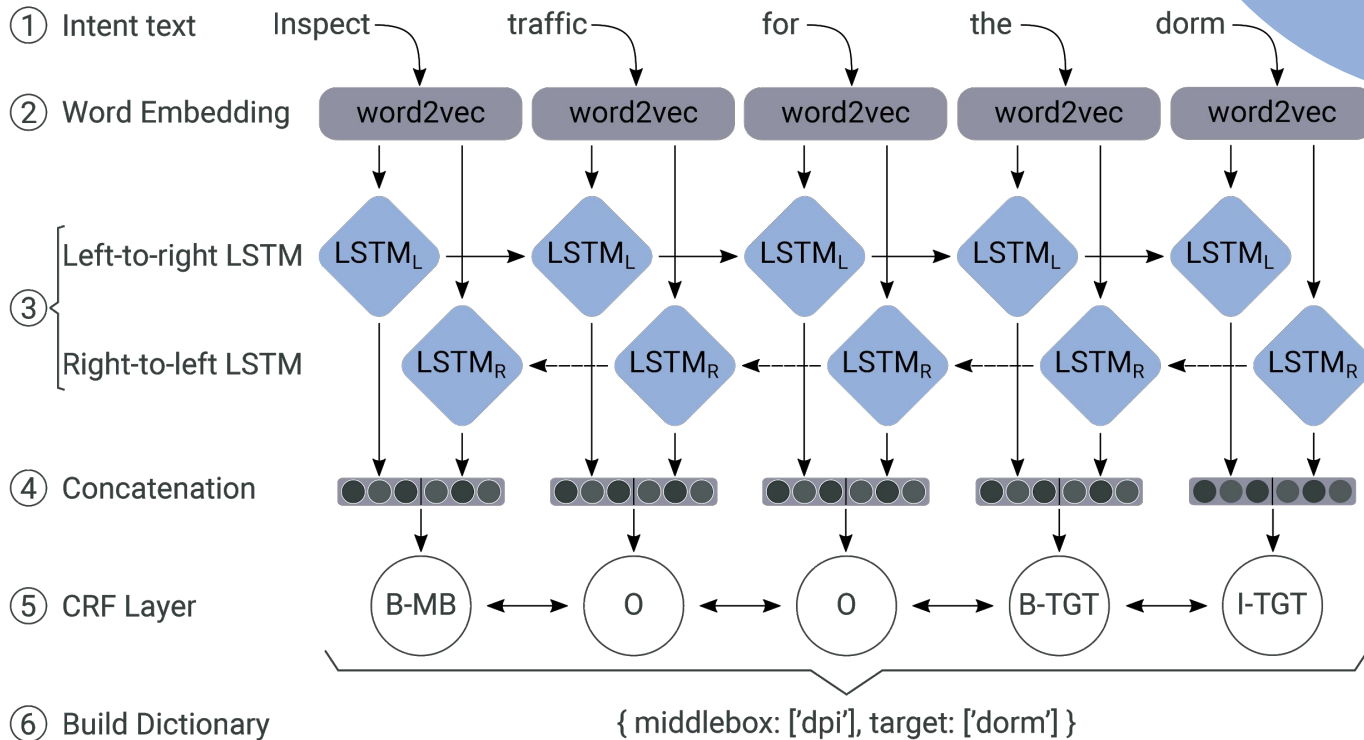


Lumi in a Nutshell

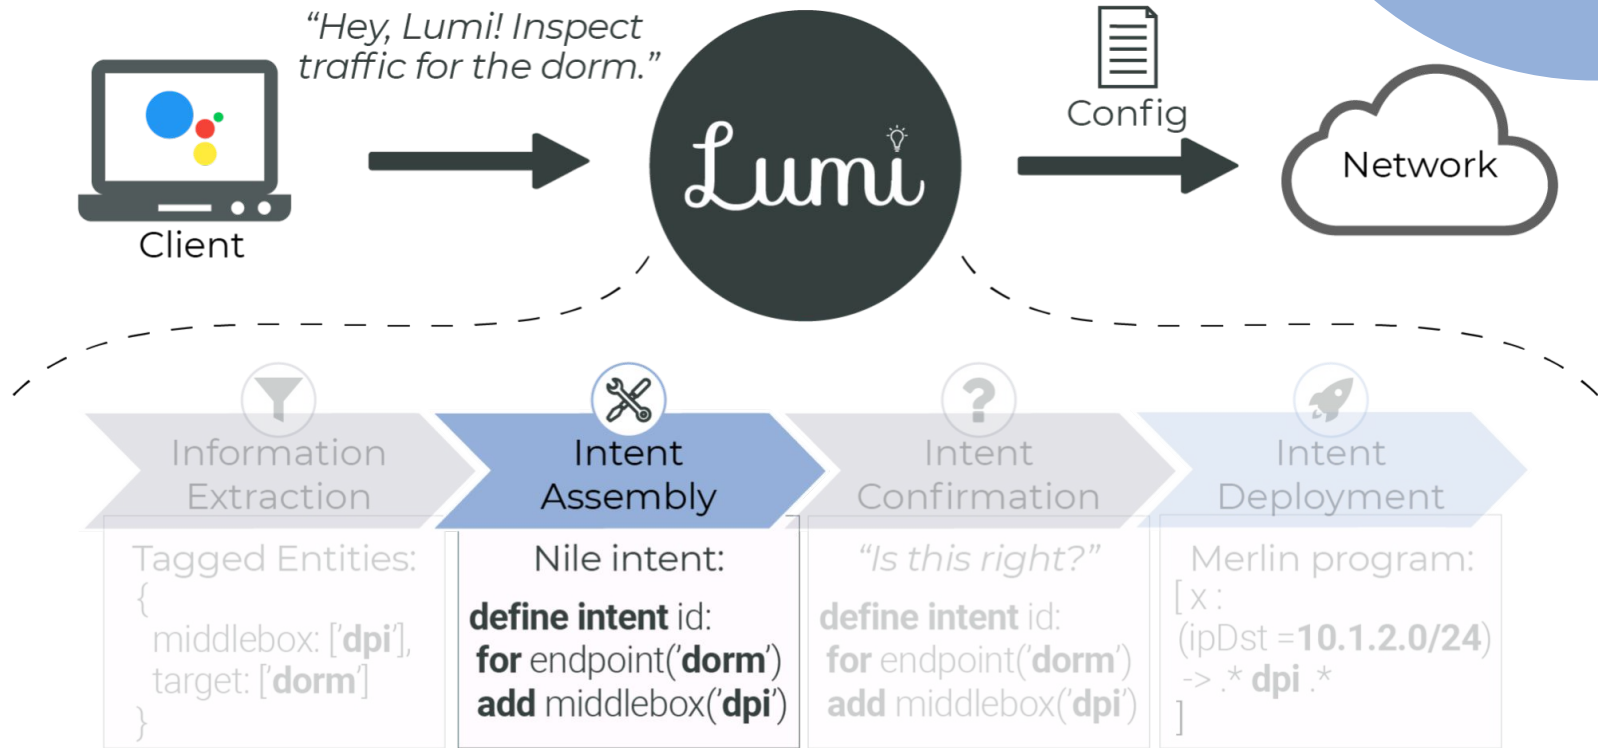


Information Extraction

Named Entity Recognition



Lumi in a Nutshell



Intent Assembly

Network Intent Language (Nile)

High legibility

High expressivity

Intent Assembly

Network Intent Language (Nile)

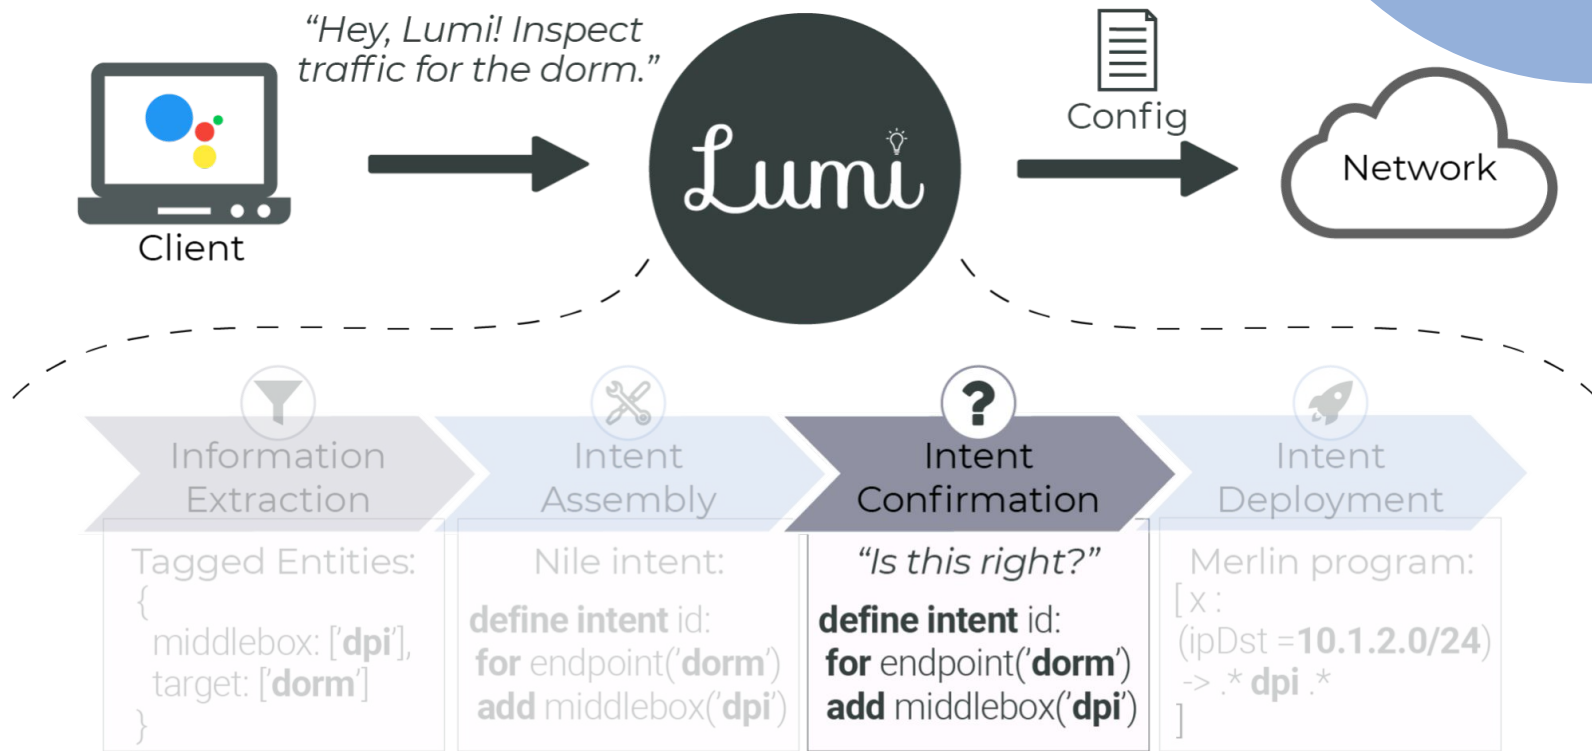
High legibility

High expressivity

“Add a firewall and intrusion detection from gateway to backend for client B with at least 100 mbps of bandwidth, and allow HTTPS only.”

```
define intent intentId:  
  from      endpoint("gateway")  
  to        endpoint("backend")  
  for       group("B")  
  add       middlebox("firewall"), middlebox("ids")  
  set       bandwidth("min", "100", "mbps")  
  allow     protocol("https")
```

Lumi in a Nutshell

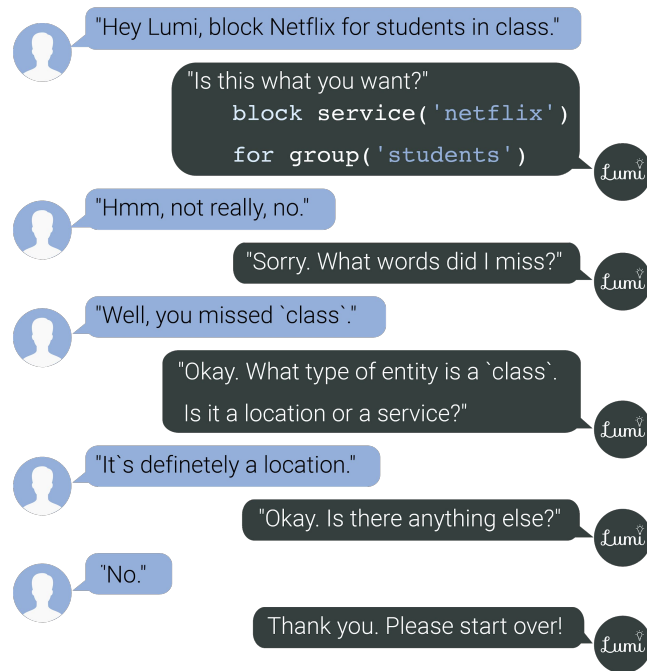


Intent Confirmation

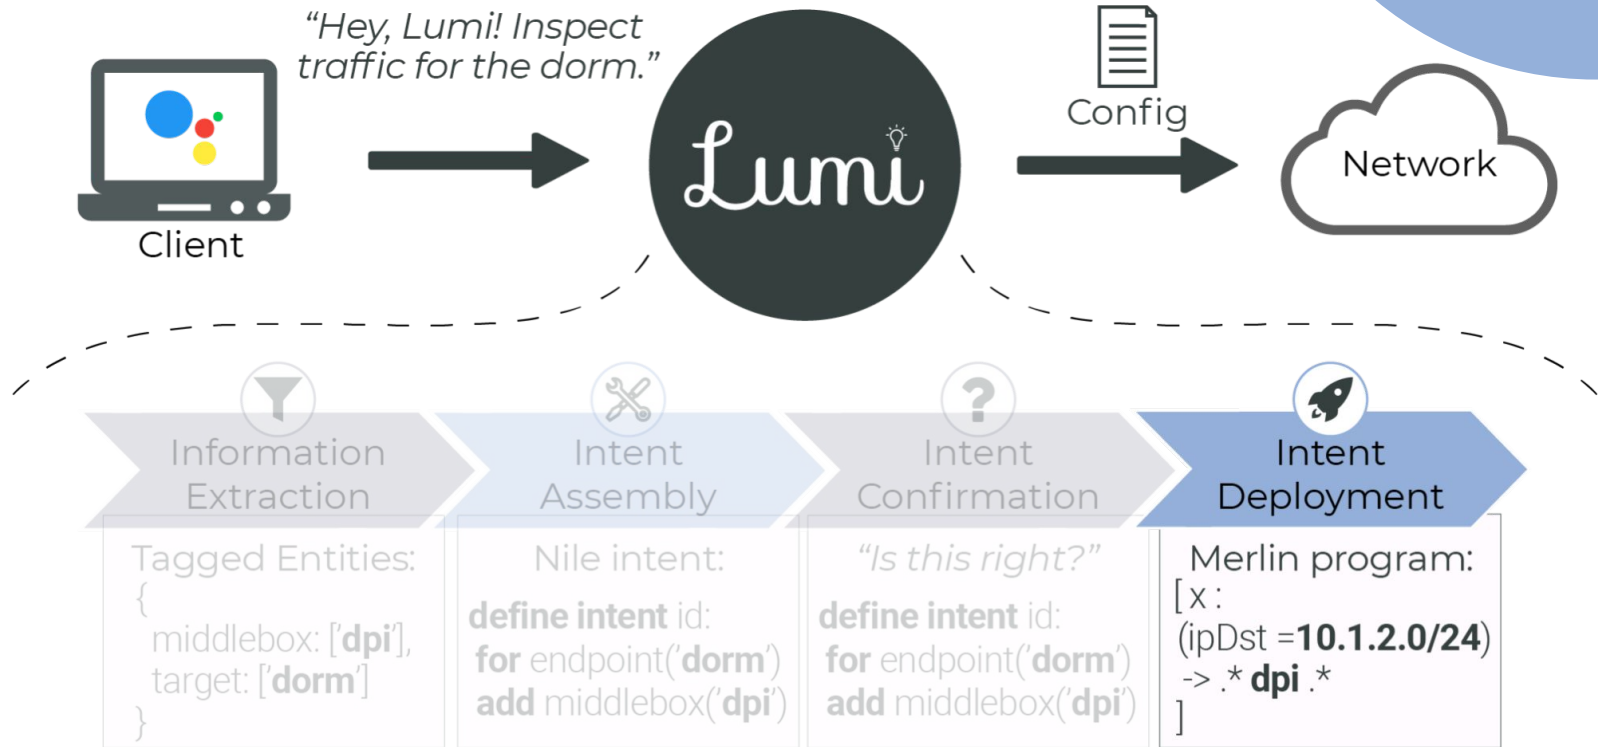
- Machine learning is always prone to mistakes
- Training data is hard to come by
 - **Use operator knowledge!**
- Iterate to extract key-value pairs
 - **Re-train NER model each time**

Intent Confirmation

- Machine learning is always prone to mistakes
- Training data is hard to come by
 - **Use operator knowledge!**
- Iterate to extract key-value pairs
 - **Re-train NER model each time**



Lumi in a Nutshell





Evaluation

Evaluation

Information Extraction

How accurate is Lumi in extracting information?

Evaluation

Information Extraction

How accurate is Lumi in extracting information?

Two datasets:

- ***alpha***: 150 Lumi-supported intents
- ***campi***: 50 intents extracted by crawling through 50 US University websites
 - Not every University website had explicit network policies

Evaluation

Campi dataset

“No individual service or system running on the wired/wireless network should use more than 10 gigabytes of bandwidth per day, regardless of whether it is in inbound or outbound over the commodity network link”

Carnegie Mellon University

“Quotas for students are 5000 Megabyte per hr download and 2000 Megabyte per hour upload”

Columbia University

Evaluation

Information Extraction

How accurate is Lumi in extracting information?

Train-test split: **75% - 25%**

Dataset	# of Entries	Precision	Recall	F1
<i>alpha</i>	150	0.966	0.987	0.991
<i>campi</i>	50	1	0.979	0.989
<i>alpha + campi</i>	200	0.992	0.969	0.980

Evaluation

Information Extraction

How accurate is Lumi in extracting information?

Train-test split: **75% - 25%**

Dataset	# of Entries	Precision	Recall	F1
<i>alpha</i>	150	0.966	0.987	0.991
<i>campi</i>	50	1	0.979	0.989
<i>alpha + campi</i>	200	0.992	0.969	0.980

Takeaway: with proper training, Lumi can accurately parse input intents.

Evaluation

Intent Confirmation and Feedback

*How much does
the feedback help?*

*How often is
feedback necessary?*

Evaluation

Intent Confirmation and Feedback

How much does the feedback help?

- Combine **alpha** and **campi** datasets
- Train with random 75% of intents
- Test with remaining 25% of intents
 - For each test intent, measure Precision and Recall
 - If there is any FP or FN, add intent to training set, and retrain model.


How often is feedback necessary?

- Repeat same experiment 30 times



Evaluation

Intent Confirmation and Feedback

How much does the feedback help?




How often is feedback necessary?




Evaluation

Intent Confirmation and Feedback

How much does the feedback help?



How often is feedback necessary?




On average, 10 out of 50 were used as feedback



Evaluation

Intent Confirmation and Feedback

How much does the feedback help?



How often is feedback necessary?



Takeaway: feedback helps Lumi learn over time, but is seldom necessary.



User Study

User Study

Completely anonymous and entirely online

- Subjects were put in the shoes of a **campus network operator**
- Asked to complete **5 tasks** using our developed Lumi prototype

The screenshot displays the Lumi chatbot interface in a browser window. The interface includes a progress bar at the top with five tasks and a 'Post questionnaire' button. The main content area shows 'Task 3: usage quotas' with a text description: 'Consider the simplified network infrastructure depicted below. Some students in the dorms have been binge-watching the new season of Game of Thrones. However, apparently, all of them have set the which is causing them to download an excessive amount of data. Please use Lumi's chatbot interface on the right to set a 10 GB per week download quota for students in dorms.'

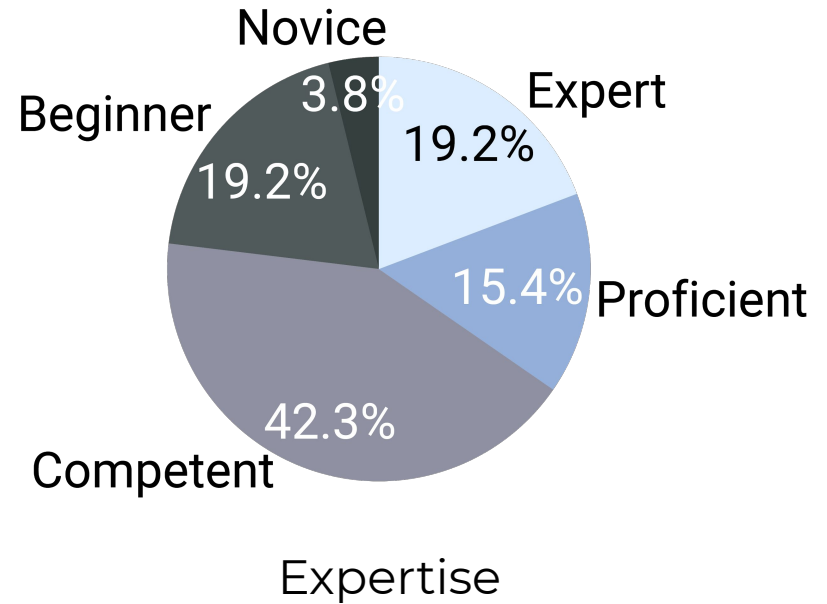

Below the text are two network diagrams. The 'Current' diagram shows a network with Internet, Gateway, DMZ, LAN, and various servers (Libs, Servers, Guests, Dorms) connected to HBD and Lumi nodes. The 'Goal' diagram shows the same network but with a '10 GB/wk' quota limit indicated on the connection to the Dorms. To the right of the diagrams is a 'Welcome to Lumi' chatbot window with a 'Type your message' input field.

Another browser window in the background shows 'Task 4: bandwidth throttling' with a similar network diagram. The text for Task 4 reads: 'Structure depicted below. Lately you have received many complaints on traffic per analysis, you realize that some services hosted in the servers have been receiving a lot of traffic. Please use Lumi's chatbot interface on the right to set a 5 Gbps bandwidth limit.' The diagram for Task 4 shows a network with Internet, Gateway, DMZ, LAN, and servers (Libs, PS, HTP, FTP, DNS, MMS) connected to HBD and Lumi nodes. A '5 Gbps From 4PM to 7PM' limit is indicated on the connection to the DMZ. A 'Welcome to Lumi' chatbot window is also visible on the right.

User Study


Subjects Profiling

- **26 Participants**




User Study

Information Extraction and Feedback



of Tasks Completed



Impact of Feedback

Summary

Lumi

- End-to-end system to allow using **natural language** for network management
- Uses **Nile** as an abstraction layer for intent confirmation.
- **Learns over time** through operator feedback.

What's next:

- Make Lumi production-ready
- Support for more features for use case scenarios other than Campus Networks

Thank you!
Questions?

Arthur Jacobs
asjacobs@inf.ufrgs.br



<http://lumichatbot.github.io>

Backup

Evaluation

Intent Deployment

How fast can Lumi compile and deploy intents?

Synthetically-generated dataset, with 30 **Nile** intents of each type:

- Middlebox chaining
- ACL
- QoS
- Temporal
- Mixed

Evaluation

Intent Deployment

How fast can Lumi compile and deploy intents?

Intent Type	Compilation Time (ms)	Deployment Time (ms)
Middlebox Chaining	4.402	110
ACL	3.115	112
Qos	3.113	132
Temporal	4.504	111
Mixed	4.621	1030

Evaluation

Intent Deployment

How fast can Lumi compile and deploy intents?


Generally, in a few milliseconds!

Intent Type	Compilation Time (ms)	Deployment Time (ms)
Middlebox Chaining	4.402	110
ACL	3.115	112
Qos	3.113	132
Temporal	4.504	111
Mixed	4.621	1030

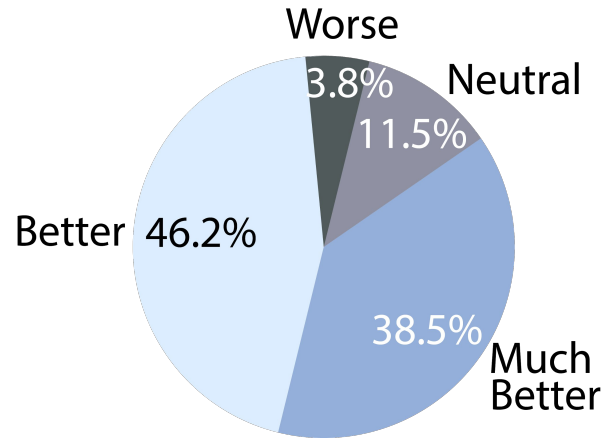
User Study

Users Reactions and Usability

How easy was Lumi to use?



Compared to traditional management, how does Lumi compare?



Would you rather use Lumi or traditional methods?

