Aprendizado de Máquina e Avaliação de Desempenho: uma dupla dinâmica engajada da teoria à prática

Edmundo de Souza e Silva[1]

Universidade Federal do Rio de Janeiro
[1]Programa de Engenharia de Sistemas e Computação, COPPE

2015

Performance Evaluation and Machine Learnig: The Dynamic Duo

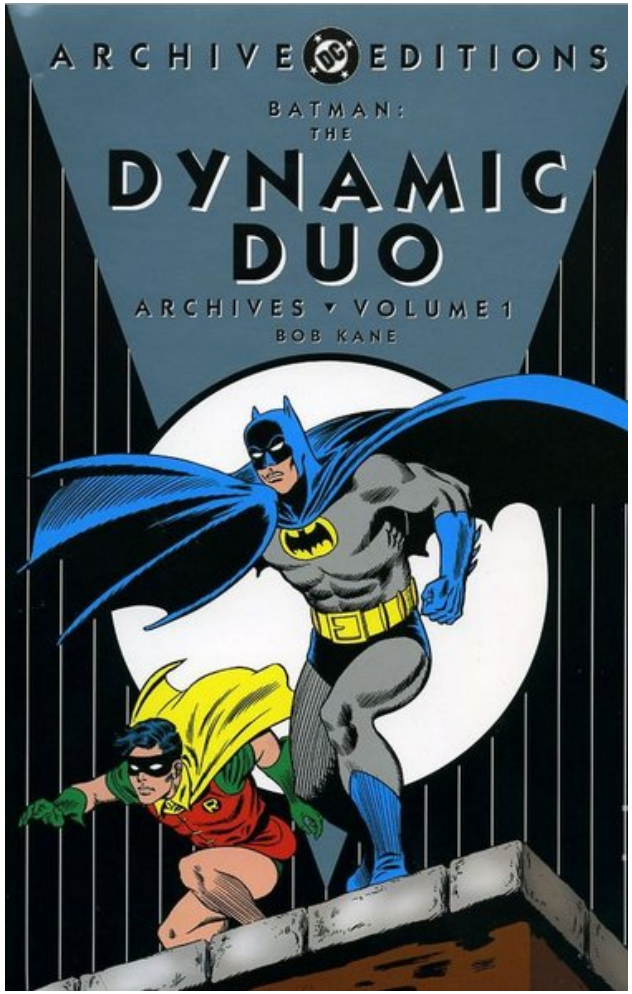

Preliminaries

- What is Computer System Modeling and Analysis?
- What is Machine Learning?
- Can we take advantage of both areas?
- What problems we address?
- Is this useful in practice?

Preliminaries

- What is Computer System Modeling and Analysis?
- What is Machine Learning?
- Can we take advantage of both areas?
- What problems we address?
- Is this useful in practice?

Preliminaries

- What is Computer System Modeling and Analysis?
- What is Machine Learning?
- Can we take advantage of both areas?
- What problems we address?
- Is this useful in practice?

Preliminaries

- What is Computer System Modeling and Analysis?
- What is Machine Learning?
- Can we take advantage of both areas?
- What problems we address?
- Is this useful in practice?

Preliminaries

- What is Computer System Modeling and Analysis?
- What is Machine Learning?
- Can we take advantage of both areas?
- What problems we address?
- Is this useful in practice?

Performance Evaluation

Performance Evaluation

- Modeling and analysis is an essential ingredient of the design process of most systems
- Devising new systems: generally needs analysis of its performance
  - What are the advantages of the new architecture?
  - which scheduling policies to use?
  - what speed to operate servers?
  - On what conditions can the system efficiently operate
- We want to predict behavior of
  - an algorithm
  - a protocol
  - a new computer architecture
  - customers accessing some system
  - ⋮
- We want to perform tradeoff analysis

Performance Evaluation

- Modeling and analysis is an essential ingredient of the design process of most systems
- Devising new systems: generally needs analysis of its performance
  - What are the advantages of the new architecture?
  - which scheduling policies to use?
  - what speed to operate servers?
  - On what conditions can the system efficiently operate
- We want to predict behavior of
  - an algorithm
  - a protocol
  - a new computer architecture
  - customers accessing some system
  - ⋮
- We want to perform tradeoff analysis

Performance Evaluation

- Modeling and analysis is an essential ingredient of the design process of most systems
- Devising new systems: generally needs analysis of its performance
  - What are the advantages of the new architecture?
  - which scheduling policies to use?
  - what speed to operate servers?
  - On what conditions can the system efficiently operate
- We want to predict behavior of
  - an algorithm
  - a protocol
  - a new computer architecture
  - customers accessing some system
  - ⋮
- We want to perform tradeoff analysis

Performance Evaluation

- Modeling and analysis is an essential ingredient of the design process of most systems
- Devising new systems: generally needs analysis of its performance
  - What are the advantages of the new architecture?
  - which scheduling policies to use?
  - what speed to operate servers?
  - On what conditions can the system efficiently operate
- We want to predict behavior of
  - an algorithm
  - a protocol
  - a new computer architecture
  - customers accessing some system
  - ⋮
- We want to perform tradeoff analysis

Examples

including research work of our group

Disk clusters

- Assume that the MTBF of a single disk is 300.000 hours. Probabilidade that a failure occurs in one disk unit during one hour $= 3.33310^{-6}$. (in one month: 0.0023942)
- Question: if you have 1000 disks, what is the probability that one disk fails in one month?
- Answer: 0.90902
- Question: if you have 100.000 disks, what is the probability that one disk fails in one hour? (common in disc clusters)
- Answer: 0.99966 $\longrightarrow$ you WILL have a disk failed somewhere in your cluster!!!

Disk clusters

- Assume that the MTBF of a single disk is 300.000 hours. Probabilidade that a failure occurs in one disk unit during one hour $= 3.33310^{-6}$. (in one month: 0.0023942)
- Question: if you have 1000 disks, what is the probability that one disk fails in one month?
- Answer: 0.90902
- Question: if you have 100.000 disks, what is the probability that one disk fails in one hour? (common in disc clusters)
- Answer: 0.99966 $\longrightarrow$ you WILL have a disk failed somewhere in your cluster!!!

Disk clusters

- Assume that the MTBF of a single disk is 300.000 hours. Probabilidade that a failure occurs in one disk unit during one hour $= 3.33310^{-6}$. (in one month: 0.0023942)
- Question: if you have 1000 disks, what is the probability that one disk fails in one month?
- Answer: 0.90902
- Question: if you have 100.000 disks, what is the probability that one disk fails in one hour? (common in disc clusters)
- Answer: 0.99966 $\longrightarrow$ you WILL have a disk failed somewhere in your cluster!!!

Disk clusters

- Assume that the MTBF of a single disk is 300.000 hours. Probabilidade that a failure occurs in one disk unit during one hour $= 3.33310^{-6}$. (in one month: 0.0023942)
- Question: if you have 1000 disks, what is the probability that one disk fails in one month?
- Answer: 0.90902
- Question: if you have 100.000 disks, what is the probability that one disk fails in one hour? (common in disc clusters)
- Answer: 0.99966 $\longrightarrow$ you WILL have a disk failed somewhere in your cluster!!!

Disk clusters

- Assume that the MTBF of a single disk is 300.000 hours. Probabilidade that a failure occurs in one disk unit during one hour $= 3.33310^{-6}$. (in one month: 0.0023942)
- Question: if you have 1000 disks, what is the probability that one disk fails in one month?
- Answer: 0.90902
- Question: if you have 100.000 disks, what is the probability that one disk fails in one hour? (common in disc clusters)
- Answer: 0.99966 $\longrightarrow$ you WILL have a disk failed somewhere in your cluster!!!

Disk clusters in Data Centers

- Data centers consume a lot of power
- Question: How can we reduce power consumption?
- Note: not all data in a large data center is accessed simultaneously $\rightarrow$ disks are not used all the time.
- Question: can we reduce power consumption by putting some disks to sleep?
- Answer: YES, but there are tradeoffs to investigate.
- What are the tradeoffs?

Disk clusters in Data Centers

- Data centers consume a lot of power
- Question: How can we reduce power consumption?
- Note: not all data in a large data center is accessed simultaneously $\rightarrow$ disks are not used all the time.
- Question: can we reduce power consumption by putting some disks to sleep?
- Answer: YES, but there are tradeoffs to investigate.
- What are the tradeoffs?

Disk clusters in Data Centers

- Data centers consume a lot of power
- Question: How can we reduce power consumption?
- Note: not all data in a large data center is accessed simultaneously $\rightarrow$ disks are not used all the time.
- Question: can we reduce power consumption by putting some disks to sleep?
- Answer: YES, but there are tradeoffs to investigate.
- What are the tradeoffs?

Disk clusters in Data Centers

- Data centers consume a lot of power
- Question: How can we reduce power consumption?
- Note: not all data in a large data center is accessed simultaneously $\rightarrow$ disks are not used all the time.
- Question: can we reduce power consumption by putting some disks to sleep?
- Answer: YES, but there are tradeoffs to investigate.
- What are the tradeoffs?

Disk clusters in Data Centers

- Data centers consume a lot of power
- Question: How can we reduce power consumption?
- Note: not all data in a large data center is accessed simultaneously $\rightarrow$ disks are not used all the time.
- Question: can we reduce power consumption by putting some disks to sleep?
- Answer: YES, but there are tradeoffs to investigate.
- What are the tradeoffs?

Disk clusters in Data Centers

- Data centers consume a lot of power
- Question: How can we reduce power consumption?
- Note: not all data in a large data center is accessed simultaneously $\rightarrow$ disks are not used all the time.
- Question: can we reduce power consumption by putting some disks to sleep?
- Answer: YES, but there are tradeoffs to investigate.
- What are the tradeoffs?

Another example

- 2 movie theaters competing with each other. They show the same movie
- Customers that arrive to see movie choose one of the theaters with equal probability (theaters are identical)
- Question: after some time:
  ( ) both theathers will receive approx. same amount of customers (and make approx. same amount of money
  ( ) one theater will get much more customers (make much more money) than the other
- Question: can we calculate the probability that one theater makes more than D dollars than the other?
- Answer: YES.

Another example

- 2 movie theaters competing with each other. They show the same movie
- Customers that arrive to see movie choose one of the theaters with equal probability (theaters are identical)
- Question: after some time:
  ( ) both theathers will receive approx. same amount of customers (and make approx. same amount of money
  ( ) one theater will get much more customers (make much more money) than the other
- Question: can we calculate the probability that one theater makes more than D dollars than the other?
- Answer: YES.

Another example

- 2 movie theaters competing with each other. They show the same movie
- Customers that arrive to see movie choose one of the theaters with equal probability (theaters are identical)
- Question: after some time:
  ( ) both theathers will receive approx. same amount of customers (and make approx. same amount of money
  ( ) one theater will get much more customers (make much more money) than the other
- Question: can we calculate the probability that one theater makes more than D dollars than the other?
- Answer: YES.

Another example

- 2 movie theaters competing with each other. They show the same movie
- Customers that arrive to see movie choose one of the theaters with equal probability (theaters are identical)
- Question: after some time:
  ( ) both theathers will receive approx. same amount of customers (and make approx. same amount of money
  ( ) one theater will get much more customers (make much more money) than the other
- Question: can we calculate the probability that one theater makes more than D dollars than the other?
- Answer: YES.

Another example

- 2 movie theaters competing with each other. They show the same movie
- Customers that arrive to see movie choose one of the theaters with equal probability (theaters are identical)
- Question: after some time:
  ( ) both theathers will receive approx. same amount of customers (and make approx. same amount of money
  ( ) one theater will get much more customers (make much more money) than the other
- Question: can we calculate the probability that one theater makes more than D dollars than the other?
- Answer: YES.

Another example

- 2 movie theaters competing with each other. They show the same movie
- Customers that arrive to see movie choose one of the theaters with equal probability (theaters are identical)
- Question: after some time:
  ( ) both theathers will receive approx. same amount of customers (and make approx. same amount of money
  ( ) one theater will get much more customers (make much more money) than the other
- Question: can we calculate the probability that one theater makes more than D dollars than the other?
- Answer: YES.
  After a long time there is a 80% chance one of the theaters got more than 20.000 customers than the other!

Modeling Cycle

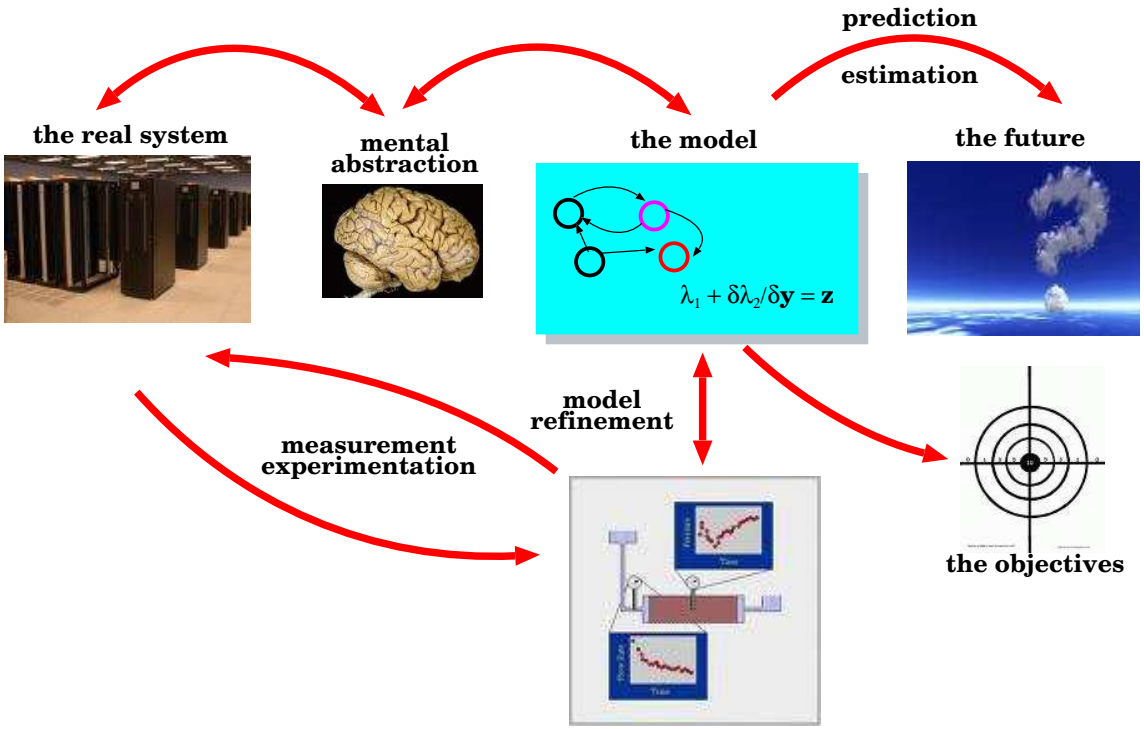

Machine Learning

Big Data

- Large amount of data produced and consumed everyday
  - social networks
  - online video streaming
  - microblogging
  - genome information
  - measurements
- How to obtain insights from data?
- What can we learn from the data?

Machine Learning

What is

Murphy:
Set is methods that can automatically detect patterns in data

Big Data

Large amount of data produced and consumed everyday

- social networks
- online video streaming
- microblogging
- genome information
- measurements

How to obtain insights from data?

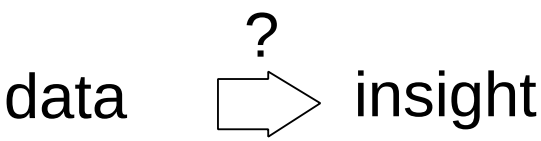

Machine Learning

What is

Murphy:
Set is methods that can automatically detect patterns in data

- Uncovered patterns $\rightarrow$:
  - predict future data
  - perform decision making
  - planning

Machine Learning

What is

Murphy:
Set is methods that can automatically detect patterns in data

- Uncovered patterns $\rightarrow$:
  - predict future data
  - perform decision making
  - planning

Obtaining insights from traces

- Machine learning = automatic pattern recognition
- Performance evaluation = model building and analysis
- machine learning tools can help to solve performance evaluation problems (and vice versa)

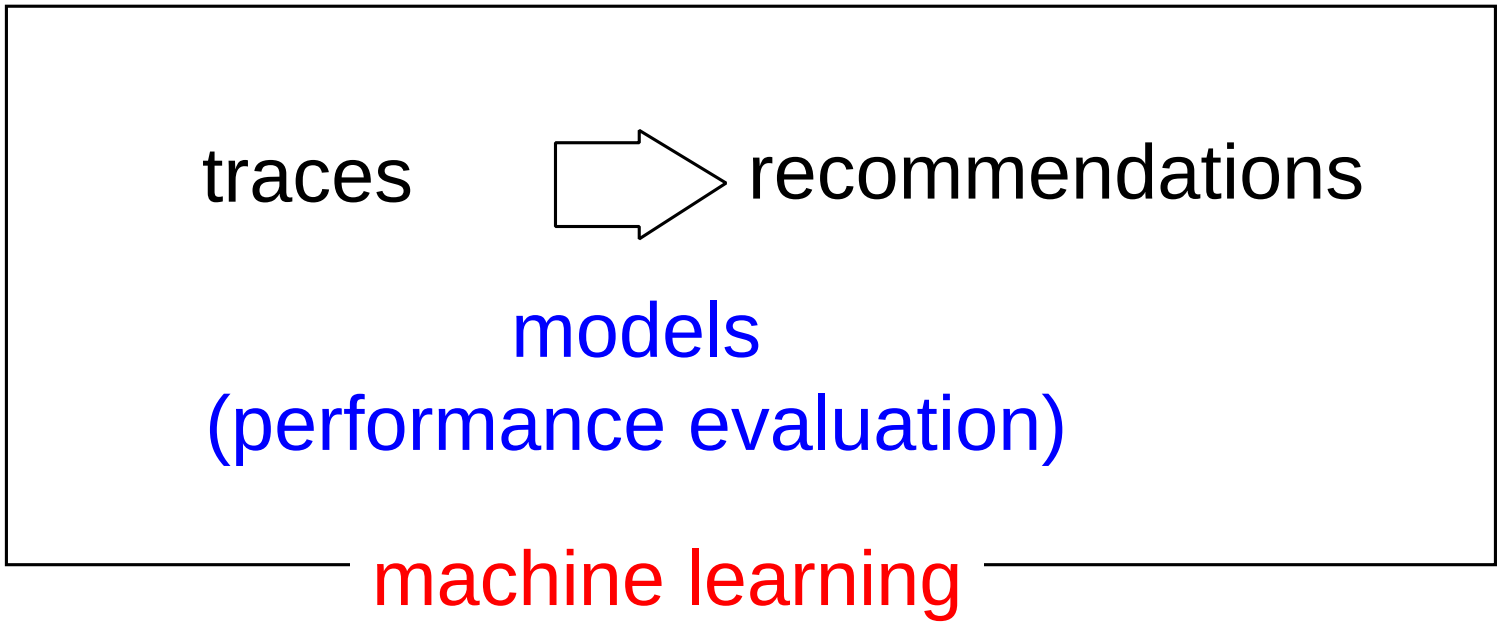

Time Series

- Given a time series, how to parameterize model to predict future values?
  - inferring customer behavior
  - modeling network channel losses
  - modeling traffic
  - generating workload
  - ⋮
- Note: we have traces of time series of one or more variables.
- Is there a structure *behind* the data?

Time Series

- Given a time series, how to parameterize model to predict future values?
  - inferring customer behavior
  - modeling network channel losses
  - modeling traffic
  - generating workload
  - ⋮
- Note: we have traces of time series of one or more variables.
- Is there a structure *behind* the data?

Time Series

- Given a time series, how to parameterize model to predict future values?
  - inferring customer behavior
  - modeling network channel losses
  - modeling traffic
  - generating workload
  - ⋮
- Note: we have traces of time series of one or more variables.
- Is there a structure *behind* the data?

Time Series

- Given a time series, how to parameterize model to predict future values?
  - inferring customer behavior
  - modeling network channel losses
  - modeling traffic
  - generating workload
  - ⋮
- Note: we have traces of time series of one or more variables.
- Is there a structure *behind* the data?

Time Series

- Given a time series, how to parameterize model to predict future values?
  - inferring customer behavior
  - modeling network channel losses
  - modeling traffic
  - generating workload
  - ⋮
- Note: we have traces of time series of one or more variables.
- Is there a structure *behind* the data?

Time Series

- Given a time series, how to parameterize model to predict future values?
  - inferring customer behavior
  - modeling network channel losses
  - modeling traffic
  - generating workload
  - ⋮
- Note: we have traces of time series of one or more variables.
- Is there a structure *behind* the data?

Time Series

- Given a time series, how to parameterize model to predict future values?
  - inferring customer behavior
  - modeling network channel losses
  - modeling traffic
  - generating workload
  - ⋮
- Note: we have traces of time series of one or more variables.
- Is there a structure *behind* the data?

Time Series

- Given a time series, how to parameterize model to predict future values?
  - inferring customer behavior
  - modeling network channel losses
  - modeling traffic
  - generating workload
  - ⋮
- Note: we have traces of time series of one or more variables.
- Is there a structure *behind* the data?

Summary

- Recall from Performance Evaluation
- Machine Learning

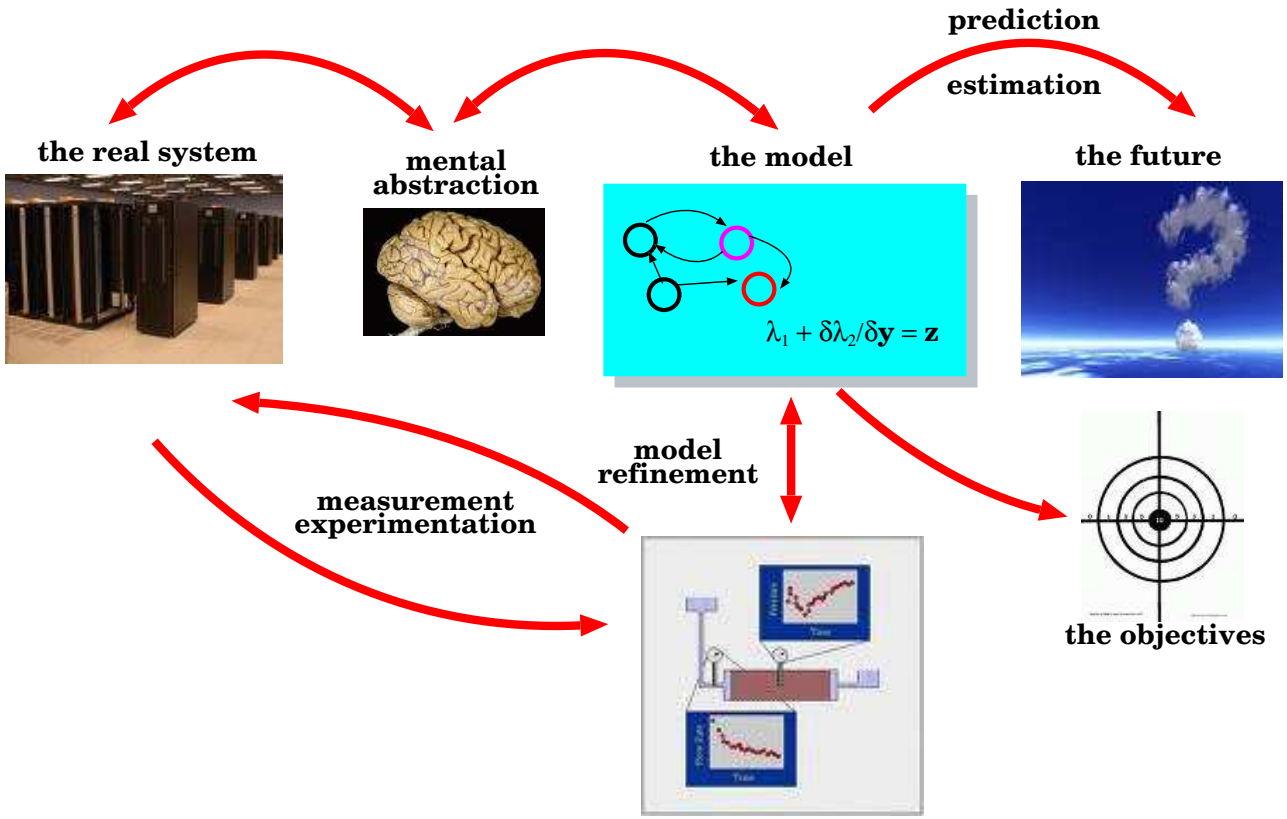

Summary

- Recall from Performance Evaluation
- Machine Learning

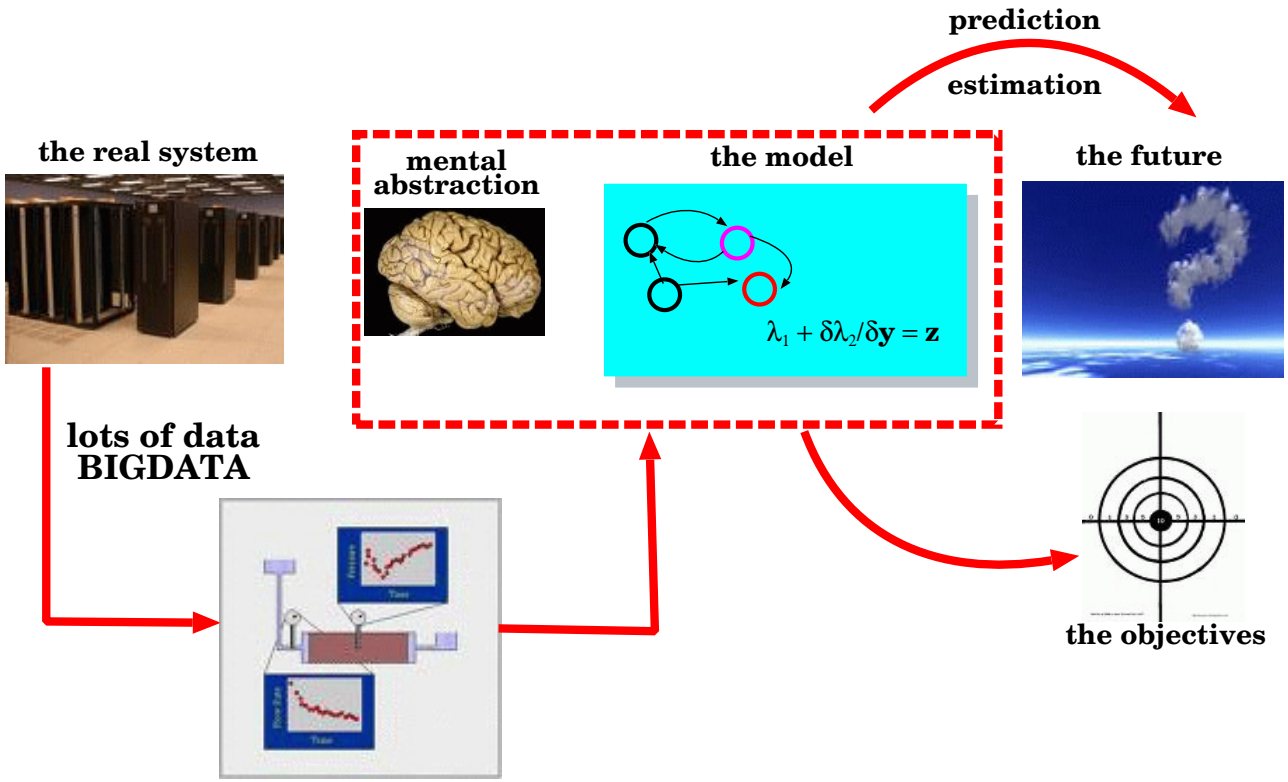

Projects of our Group

Distance Learning Initiative

P&D - service

videoaula@RNP

Distance Learning Initiative

- More than 800 videolectures (approximately 40-90 minutes each)
- Technology completely developed at the university
- CEDERJ Computer Science course started in 2005
- It has been a **service of the RNP** since 2011: Videoaula@RNP
- Designs started as a research project (CNPq - FAPERJ projects) and later made into a product (supported by RNP) and transfered to RNP.

Distance Learning Initiative

- More than 800 videolectures (approximately 40-90 minutes each)
- Technology completely developed at the university
- CEDERJ Computer Science course started in 2005
- It has been a **service of the RNP** since 2011: Videoaula@RNP
- Designs started as a research project (CNPq - FAPERJ projects) and later made into a product (supported by RNP) and transfered to RNP.

Distance Learning Initiative

- More than 800 videolectures (approximately 40-90 minutes each)
- Technology completely developed at the university
- CEDERJ Computer Science course started in 2005
- It has been a **service of the RNP** since 2011: Videoaula@RNP
- Designs started as a research project (CNPq - FAPERJ projects) and later made into a product (supported by RNP) and transfered to RNP.

Distance Learning Initiative

- More than 800 videolectures (approximately 40-90 minutes each)
- Technology completely developed at the university
- CEDERJ Computer Science course started in 2005
- It has been a **service of the RNP** since 2011: Videoaula@RNP
- Designs started as a research project (CNPq - FAPERJ projects) and later made into a product (supported by RNP) and transfered to RNP.

The Service videoaula@RNP

Example

The Service videoaula@RNP

Example

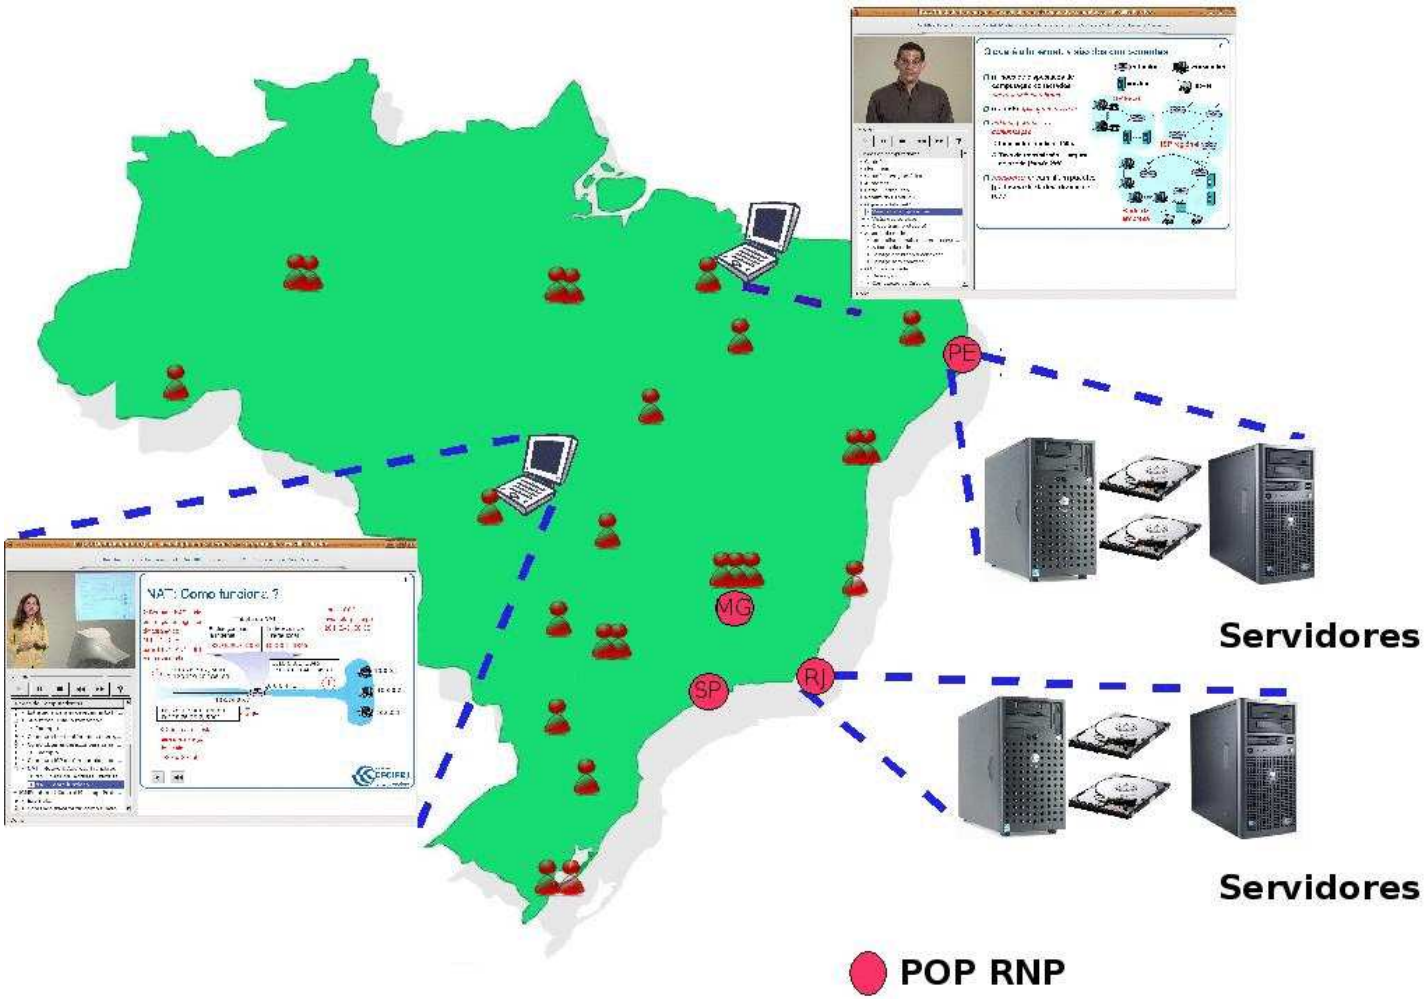

Videoaula@RNP Service
Usage

- Daily access: almost 7,000 accesses in one day (Feb)
- Reached more than 110,000 accesses in one month

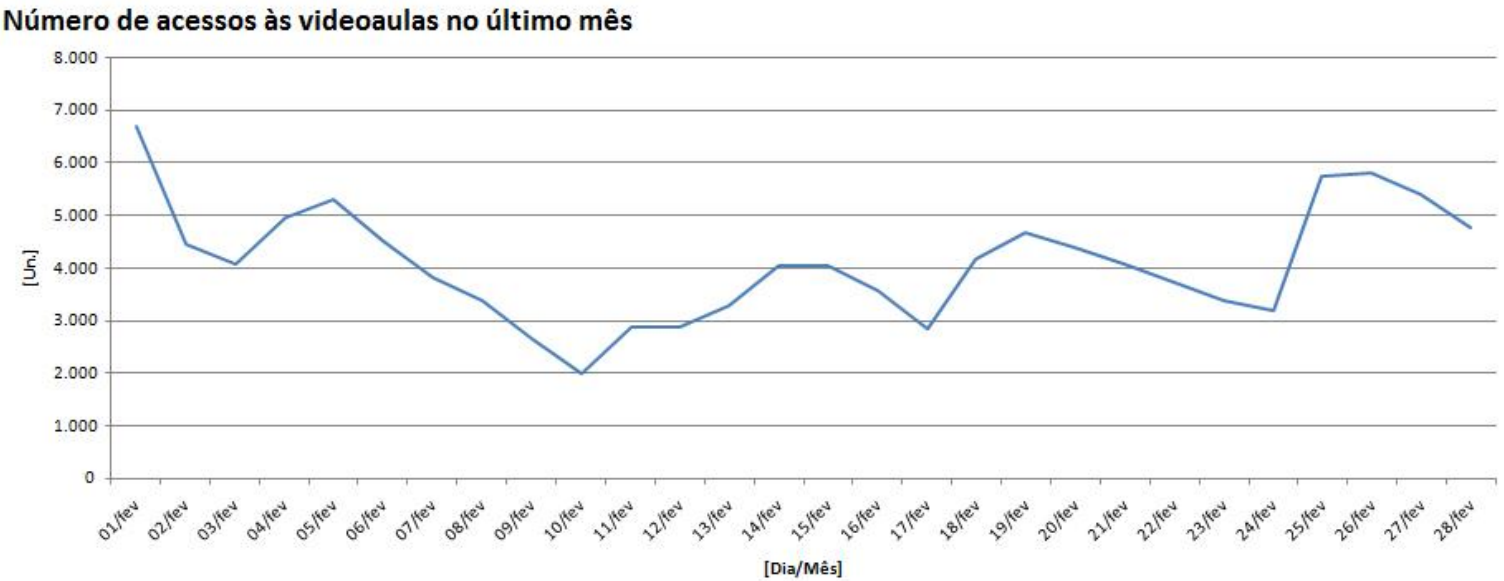

Project supported by GOOGLE

Project

An Intelligent Recommendation System based on Video Lectures for Distance Education

Project supported by Google

ESTADO DE S. PAULO

Economia | B15

Google financia pesquisas sobre web no Brasil

Empresa investe em estudos acadêmicos que tentam desvendar do acesso de vídeos virais ao comportamento do aluno de videoaulas

Empresa investe em estudos académicos do...
... ao comportamento de alunos de videoaulas

Project supported by Google

Project supported by Google

- Reportagem Folha do estado de São Paulo, 10 de Junho de 2015
- http://www1.folha.uol.com.br/tec/2015/06/1633848-google-dara-bolsas-de-mestrado-e-doutorado-em-computacao-no-brasil.shtml

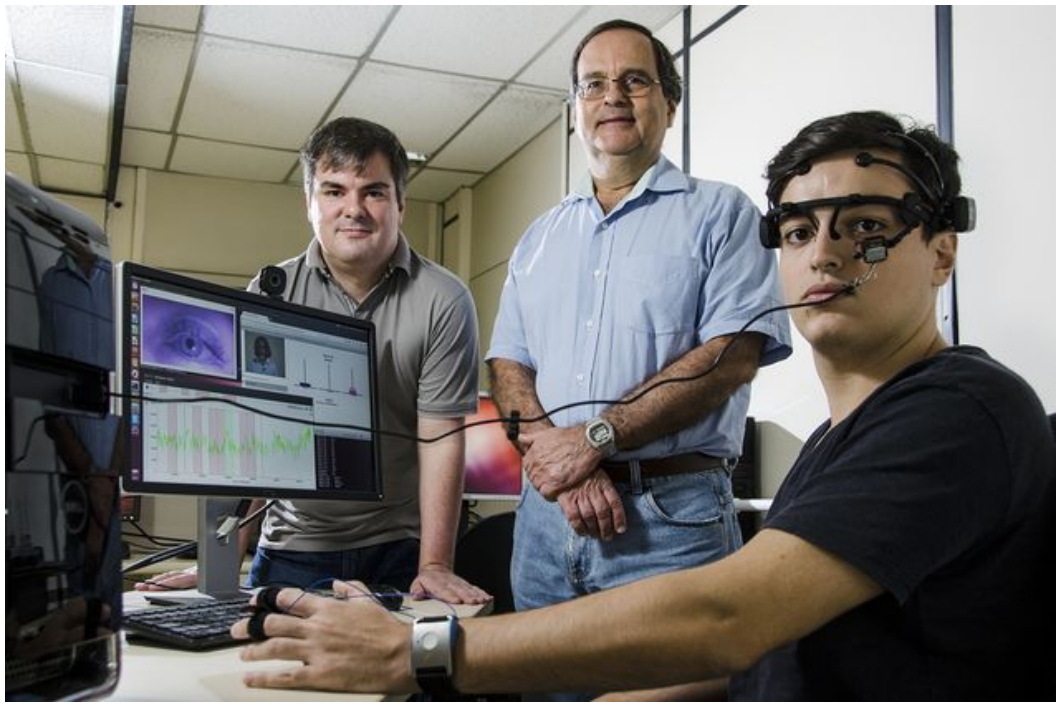

Project: Recommendation System for Videolectures

Our Objectives

- Develop an Intelligent Recommendation System based on VideoLectures for Distance Education
- Research Goals:
  - To adapt the videolecture material according to individual user's needs.
  - Automatically make suggestions to each student in realtime:
    - additional written material on the subject
    - additional explanation from short videos
    - additional exercises
  - Give important feedback to faculty involved on each class

Project: Recommendation System for Videolectures

Our Objectives

- Develop an Intelligent Recommendation System based on VideoLectures for Distance Education
- Research Goals:
  - To adapt the videolecture material according to individual user's needs.
  - Automatically make suggestions to each student in realtime:
    - additional written material on the subject
    - additional explanation from short videos
    - additional exercises
  - Give important feedback to faculty involved on each class

Project: Recommendation System for Videolectures

Our Objectives

- Develop an Intelligent Recommendation System based on VideoLectures for Distance Education
- Research Goals:
  - To adapt the videolecture material according to individual user's needs.
  - Automatically make suggestions to each student in realtime:
    - additional written material on the subject
    - additional explanation from short videos
    - additional exercises
  - Give important feedback to faculty involved on each class

Project: Recommendation System for Videolectures

Our Objectives

- Develop an Intelligent Recommendation System based on VideoLectures for Distance Education
- Research Goals:
  - To adapt the videolecture material according to individual user's needs.
  - Automatically make suggestions to each student in realtime:
    - additional written material on the subject
    - additional explanation from short videos
    - additional exercises
  - Give important feedback to faculty involved on each class

Project: Recommendation System for Videolectures

Our Objectives

- Develop an Intelligent Recommendation System based on VideoLectures for Distance Education
- Research Goals:
  - To adapt the videolecture material according to individual user's needs.
  - Automatically make suggestions to each student in realtime:
    - additional written material on the subject
    - additional explanation from short videos
    - additional exercises
  - Give important feedback to faculty involved on each class

Project: Recommendation System for Videolectures

Our Objectives

- Develop an Intelligent Recommendation System based on VideoLectures for Distance Education
- Research Goals:
  - To adapt the videolecture material according to individual user's needs.
  - Automatically make suggestions to each student in realtime:
    - additional written material on the subject
    - additional explanation from short videos
    - additional exercises
  - Give important feedback to faculty involved on each class

The System

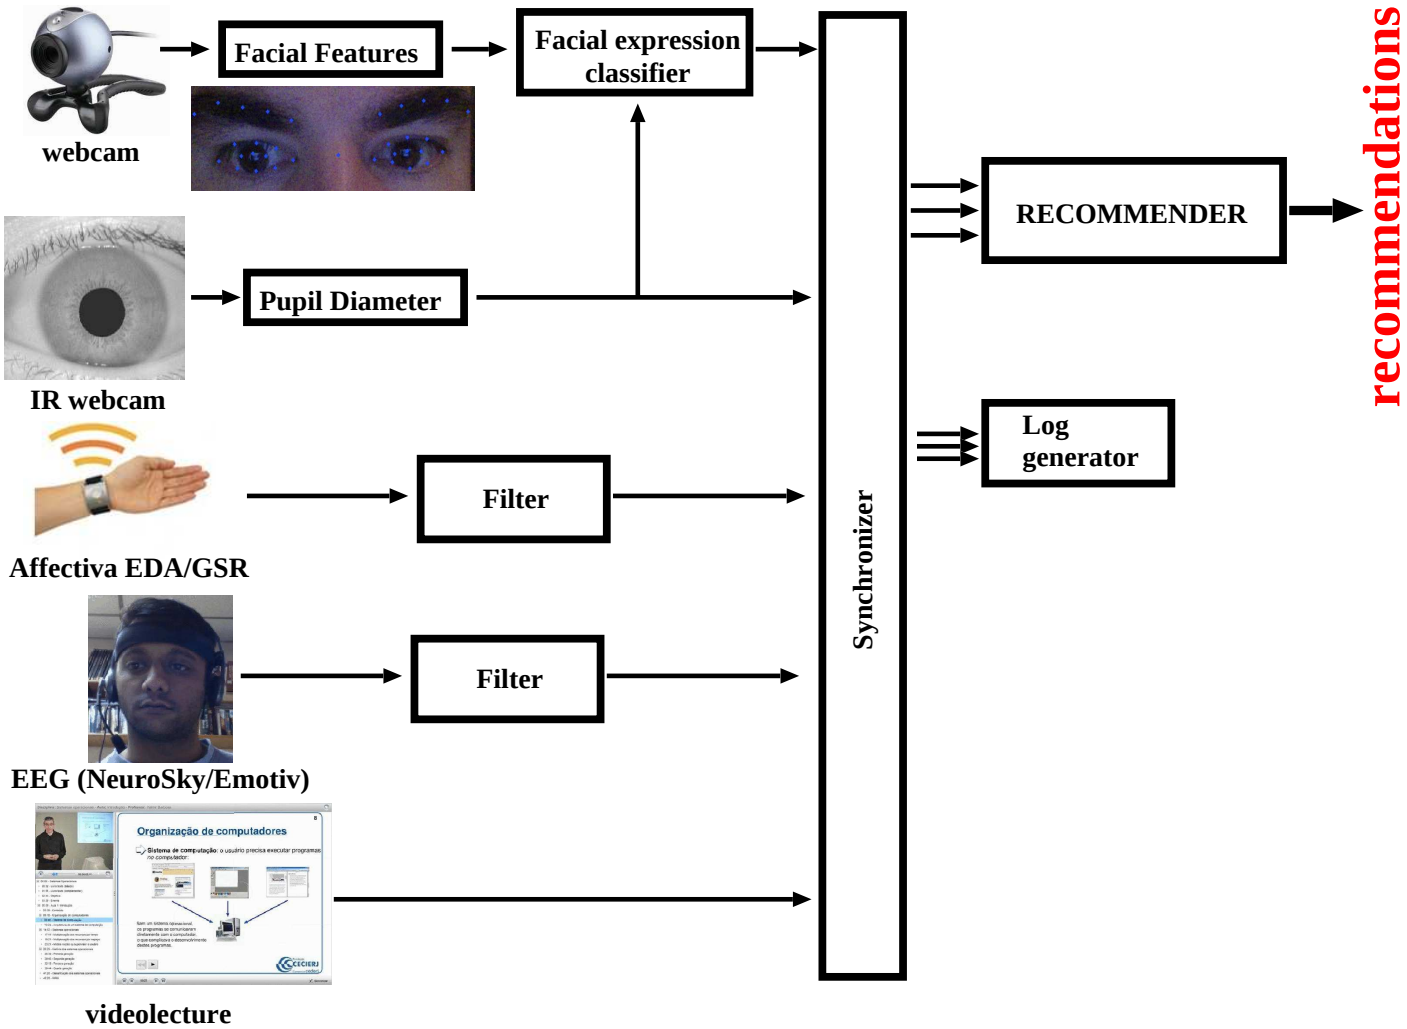
webcam
Facial Features
Facial expression classifier
IR webcam
Pupil Diameter
Affectiva EDA/GSR
Filter
EEG (NeuroSky/Emotiv)
Filter
Organização de computadores
videolecture
Synchronizer
RECOMMENDER
Log generator
recommendations

Startup - Measurements and Planning

TGR - Tecnologia em Gestão e Planejamento de Redes

Measurements and Analysis

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
  Communications of the ACM - November 2012
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
  Impact of Home Networks on End-to-End Performance (HomNets 2010)
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
  Broadband Internet Performance: A View From the Gateway (ACM/SIGCOMM 2011)
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
  Speed Measurements of Residential Internet Access (IEEE/PAM 2012)
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
  Measuring Home Broadband Performance
  (Communications of the ACM - November 2012)
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
  The Realities of Home Broadband (CACM Nov 2012)
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
  Mixture Models of Endhost Network Traffic (Infocom 2013)
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
  Predicting user dissatisfaction with Internet application performance at end-hosts (Infocom 2013)
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Parceria Universidade/Empresa

- Assunto científico de interesse internacional:
- Artigos recentes têm sido publicados em veículos internacionais de renome (2012, 2013)
- "much remains to be done to improve our understanding of broadband services"
- Transferência de novas técnicas desenvolvidas no LAND para TGR
- Desenvolvimento teórico aliado a experimentação em campo
- Ganhos para a sociedade: tópico de interesse para formuladores de políticas públicas e consumidores
- Pesquisa de ponta: Tese de mestrado 2015: Caracterização e modelos para avaliar o desempenho de redes de acesso residencial baseados em aprendizado de máquina .

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga.
- Software para diagnosticar problemas de instalação e manutenção
- Software para dimensionamento da capacidade da rede
- Conhecer o tráfego do cliente
- Entender o comportamento do usuário de banda larga

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga. Garantia de Qualidade $\rightarrow$ economia de recursos
- Software para diagnosticar problemas de instalação e manutenção
- Software para dimensionamento da capacidade da rede
- Conhecer o tráfego do cliente
- Entender o comportamento do usuário de banda larga

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga.
- Software para diagnosticar problemas de instalação e manutenção
- Software para dimensionamento da capacidade da rede
- Conhecer o tráfego do cliente
- Entender o comportamento do usuário de banda larga

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga.
- Software para diagnosticar problemas de instalação e manutenção Redução de custos
- Software para dimensionamento da capacidade da rede
- Conhecer o tráfego do cliente
- Entender o comportamento do usuário de banda larga

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga.
- Software para diagnosticar problemas de instalação e manutenção
- Software para dimensionamento da capacidade da rede
- Conhecer o tráfego do cliente
- Entender o comportamento do usuário de banda larga

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga.
- Software para diagnosticar problemas de instalação e manutenção
- Software para dimensionamento da capacidade da rede
  Planejamento futuro $\rightarrow$ **economia de recursos**
- Conhecer o tráfego do cliente
- Entender o comportamento do usuário de banda larga

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga.
- Software para diagnosticar problemas de instalação e manutenção
- Software para dimensionamento da capacidade da rede
- Conhecer o tráfego do cliente
- Entender o comportamento do usuário de banda larga

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga.
- Software para diagnosticar problemas de instalação e manutenção
- Software para dimensionamento da capacidade da rede
- Conhecer o tráfego do cliente Planejamento $\rightarrow$ **economia de recursos**
- Entender o comportamento do usuário de banda larga

Produtos Iniciais

- Ferramenta de medição e diagnóstico da qualidade da banda larga.
- Software para diagnosticar problemas de instalação e manutenção
- Software para dimensionamento da capacidade da rede
- Conhecer o tráfego do cliente
- Entender o comportamento do usuário de banda larga

Projeto

Implementação do software no firmware da INTEL (gateway embedded solution)

- INTEL
- Implementação do software no firmware em escala global

Gateway Inteligente

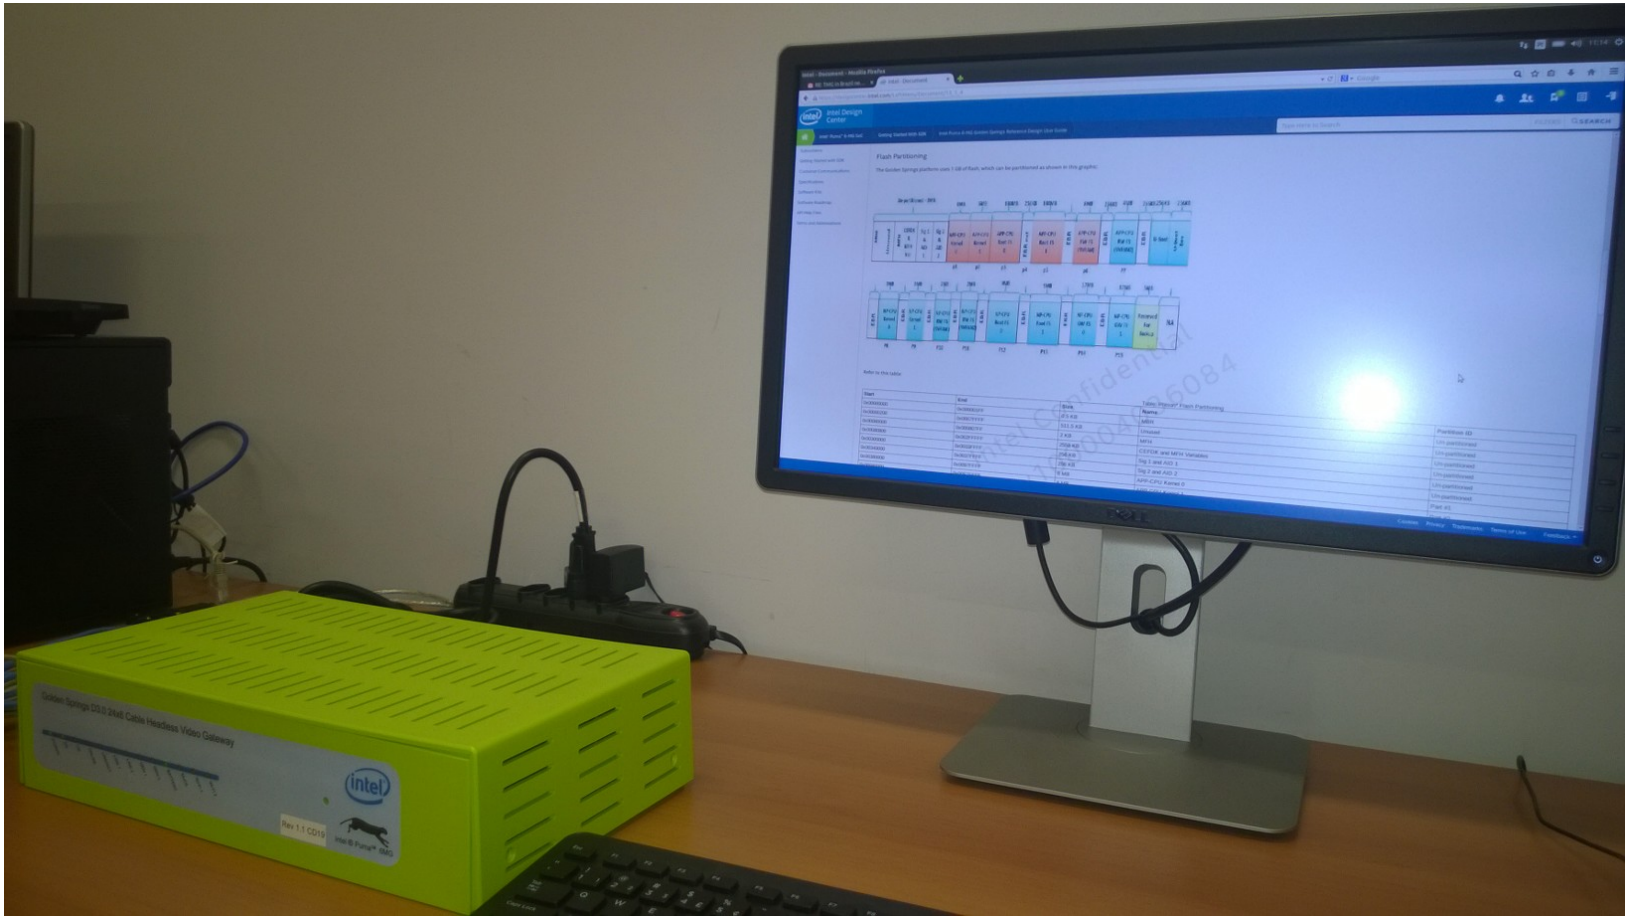

Research/Development in our Group

- Fault Tolerance is Essencial
- Performance always matter

  Stuart Feldman

  ACM's Software System Award, Vice-President Eng.

  Google

Research/Development in our Group

- Fault Tolerance is Essencial
- Performance always matter

  Stuart Feldman

  ACM's Software System Award, Vice-President Eng.

  Google

Research/Development in our Group

- Fault Tolerance is Essencial
- Performance always matter
  Stuart Feldman
  ACM's Software System Award, Vice-President Eng.
  Google

FIM

OBRIGADO!

PERGUNTAS?

www.abc.org.br/~edmundo

www.land.ufrj.br/~edmundo