# Portfolio MSc Projects

Bayesian Intelligent Autonomous Systems lab

February 19, 2024

# 1 Cocktail party problem

## Problem description

You are challenged to design an agent that learns to solve the cocktail party problem through on-the-spot interactions with a (human) listener. The cocktail party problem refers to the issue of not being able to understand your conversation partner in the presence of many simultaneously competing voices (Fig.1). The learning protocol is displayed in Fig.2. A listener wears earbuds that are capable to process audio signals in real-time (like hearing aids). In response to a detected problem, the agent proposes the most promising alternative parameter settings for the audio algorithm (the TRY step). Next, the new audio algorithm is executed in the ear buds and evaluated by the listener (EXECUTE and EVALUATE steps). Based on the listener's appraisal, the agent should now update its model of the world (LEARN step). This design loop repeats in real-time until the listener indicates that the problem has been solved.

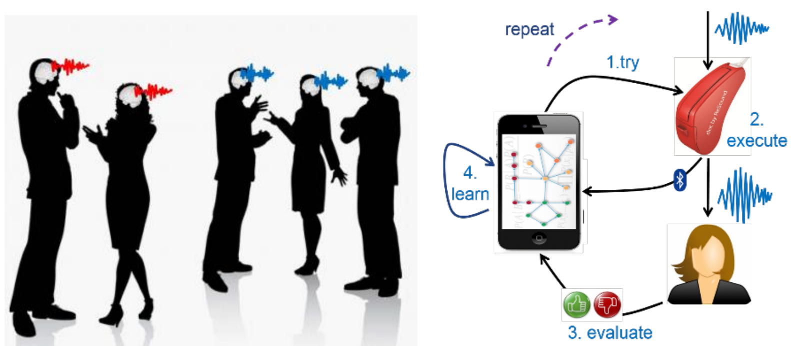


Figure 1: (Left) Cocktail party. (Right) Hearing aid.

## Student task description

This project will get you involved with the latest artificial intelligence methods, since the agent needs to (1) learn from each interaction and (2) be smart about selecting the most promising algorithm candidates. It will also give you an opportunity to learn about how biological brains solve real-time design issues.

**Concrete tasks:**

- Familiarize yourself with the literature on Active inference.
- Discuss problem and solution proposal with researchers at BIASlab.
- Implement solution proposal using BIASlab resources.
- Experiment with the implementation and compare to current solutions.

- Analyze results and reflect on what has been achieved.
- Write a report detailing the advantages and limitations of this approach.

## Project positioning

The Bayesian Intelligent Autonomous Systems lab BIASlab (http://biaslab.org, FLUX-7.060) is a subgroup of the Signal Processing Systems (SPS) that aims to develop Intelligent Autonomous Agents. These agents interact with their environment through their sensors and actuators in order to learn purposeful behavior, e.g., to navigate, play soccer or they may learn to decode speech signals under bad acoustic conditions. Our research projects are inspired by the latest insights from machine learning, computational neuroscience and signal processing.

**Supervision:**

- Weekly progress meetings with Bert de Vries.
- Periodic progress meetings with contacts at GN Resound.
- The student should prepare update meetings, preferably with derivations or visualizations in interactive notebooks.

## Timeline

The project is available from September 2017 onwards. The total duration will be 32 weeks. Halfway through, there will be a ”midterm” evaluation where the student must report on their activities and indicate how they will proceed for the remainder of the project. At the end, the student will write a thesis in 10-page double-column IEEE Transactions style summarizing their work, their findings and possible future steps. The thesis will be presented in an official ”defense” ceremony and a committee of experts will grade the student’s work.

# 2 Probabilistic Programming for Printhead Diagnostics

## Problem description

Industrial printers output thousands of pages per day. This enormous throughput is realized by a delicate printhead that operates under extremely varying conditions. The printhead consists of a few thousand micronozzles that jet ink to the paper. Through test pages, nozzles are diagnosed and when too many nozzles are failing, a maintenance action is triggered.

With the enormous throughput of industrial printers, it becomes paramount to keep printers operational. Any issues should be quickly diagnosed and resolved. Therefore, it helps if maintenance engineers are informed about potential root causes of a problem. This will help in understanding the optimal maintenance action, such as a cleaning versus a printhead replacement.

A probabilistic dynamical model of the printhead operation, based on expert insights, can consider uncertainties on the failure causes and the time-varying operational conditions. With probabilistic programming techniques, a root cause of failure can be automatically inferred from the probabilistic model, and communicated to the maintenance engineer.

## Student task description

The goal of this project is to apply probabilistic programming techniques to model printhead nozzle (mal)function, and automatically identify root causes of failure through probabilistic inference.

#### Concrete tasks:

1. Study models of nozzle malfunction;
2. Study probabilistic programming;
3. Build a probabilistic dynamical model of printhead nozzle (mal)function;
4. Given a dataset of nozzle diagnostics over time, infer states and parameters in the model;
5. Repeat 3-5 until convergence or until time is up.

## Project positioning

This project is a collaboration between Canon production printing, TNO-ESI and BIASlab. Canon Production Printing (CPP) works on game-changing technology that shapes the way that people work and communicate. TNO-ESI (esi.nl) develops new methods and techniques for system design and engineering for the high-tech industry. BIASlab (biaslab.org) works on probabilistic programming for intelligent agents that operate under real-time, situated conditions.

# 3 Active inference for robot locomotion

## Problem description

The goal is to develop an intelligent autonomous system (agent) for a quadrupedal robot (see example in Fig. 2). The agent must learn to walk: it will have to gradually build a locomotion model from interaction with its environment. You will use *Active Inference* (AIF), a probabilistic machine learning framework from the computational neuroscience community, to design and train the agent.

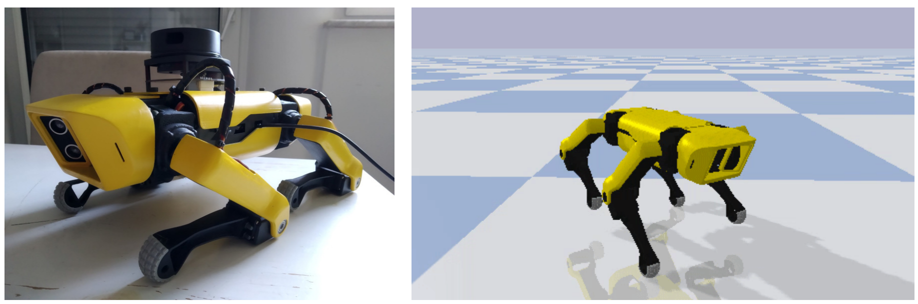

Figure 2: (Left) A physical SpotMicro. (Right) A SpotMicro in simulation. Figures courtesy of `https://spotmicroai.readthedocs.io/en/latest/`.

Please note that this is a software project; it does not involve hardware. You will write code to control a simulated robot that interacts with a simulated environment (see Fig. 2 right), generated using the open-source physics engine Bullet. The challenge will be to adapt an existing active inference agent to a specific walker.

## Project positioning

This project does not involve a company. You will be working in the Bayesian Intelligent Autonomous Systems lab (BIASlab) at TU Eindhoven. BIASlab develops probabilistic models and inference algorithms for signal processing and control problems. This graduation project is a part of a larger research initiative towards AIF agents for robot locomotion.

## Student assignment

You will initially be spending time familiarizing yourself with the tools and techniques that BIASlab develops (RxInfer.jl). Once familiar, you will write your own active inference agent based off of existing AIF agent implementations within BIASlab. Note that you'll be supported by BIASlab researchers that are working on other robot locomotion projects. Concretely, tasks include:

- Review literature on AIF agents for robotics.

- Familiarize yourself with the challenges of quadrupedal locomotion.
- Learn to use the probabilistic machine learning toolbox RxInfer.jl.
- Familiarize yourself with the simulation environment Bullet.
- Develop an active inference agent for a quadrupedal robot system.
- Reflect on what has been achieved and discuss with BIASlab's researchers.
- Write a report detailing your agent's properties and behaviour.

## Supervision

- Weekly progress meetings with dr. Wouter Kouw.
- Periodic progress meetings with BIASlab researchers.
- The student should prepare for meetings in advance, e.g. with notebooks.
- All software that has been developed should be accessible online through BIASlab's github organization and should be legible / usable for future students.

## Timeline

The project is available from September 2022 onwards. The total duration will be 32 weeks. Halfway through, there will be a "midterm" evaluation where the student must report on their activities and indicate how they will proceed for the remainder of the project. At the end, the student will write a thesis summarizing their work, their findings and possible future steps. The thesis will be presented in an official "defense" ceremony and a committee of experts will grade the student's work.

# 4 Modeling resilience of elderly patients after cancer treatment

## Problem description

Success of cancer treatment in elderly patients depends on their ability to "bounce back" to good health after treatment. Measures of this so-called resilience allow for personalized care during recovery, and may greatly impact the quality of life after cancer treatment.

More precisely, resilience is the ability of a patient to recover their health after a health stressor (e.g. a surgical intervention or chemotherapy). Dynamic indicators are physiologic and cognitive responses of a patient, measured over time. Patterns between these dynamic indicators are indicative of resilience.

During this project you will develop a probabilistic model of dynamic indicators of resilience. From measurements, probabilistic inference on this model then predicts a patient's resilience, which may support personalized decisions towards recovery.

## Student task description

Probabilistic programming is a relatively new field of research that automates the process of probabilistic inference. The goal of this project is to apply probabilistic programming techniques to model dynamic indicators of resilience in elderly patients that are recovering from cancer treatment.

**Concrete tasks:**

1. Study dynamic indicators of resilience;
2. Study probabilistic programming;
3. Build a probabilistic dynamical model of indicators of resilience;
4. Given a set of measurements, predict personalized resilience over time;
5. Repeat 3-5 until convergence or until time is up.

## Project positioning

This project is a collaboration between Orikami and BIASlab. Orikami (orikami.ai) enables personalized health care by combining medical knowledge with data science through an academic approach. BIASlab (biaslab.org) studies intelligent systems that support decision-making under situated, dynamical conditions.

# 5 Dynamic Modeling of Ink Curing and Adhesion

## Problem description

Ink for production printing applications hardens through polymerization. Under ultraviolet light exposure, long molecular chains are formed that stabilize and bind the ink to a medium. Ink adhesion quality then depends on a curing profile of the ink film, in depth and over time.

Dynamic models for ink curing are under development. However, the formal relationship between ink curing and adhesion is still unknown. Moreover, ultraviolet light exposure over the medium is never uniform, affecting the curing process in intricate ways. Therefore, uncertainty is present in many forms. Namely, in light exposure levels, the formal connection between curing and adhesion, and in the dynamic model of ink curing itself.

The scarce available knowledge should be exploited in modeling the phenomena that underlie ink adhesion. Therefore, you will use the probabilistic modeling approach, which allows for incorporating uncertainty and prior knowledge.

## Student task description

In this project you will build and train dynamic generative models of ink curing and relate these models to measures of good and bad ink adhesion.

**Concrete tasks:**

1. Study the physical chemistry of ink polymerization;
2. Study the probabilistic generative modeling approach;
3. Using available knowledge, build and train a generative dynamic model of ink curing;
4. Differentiate your model on examples of good and bad ink adhesion;
5. Use your trained models in a model comparison setting to predict ink adhesion quality from a dynamic curing profile.

## Project positioning

Canon production printing (cpp.canon) works on game-changing technology that shapes the way that people work and communicate. BIASlab (biaslab.org) works in neuroscience-inspired artificial intelligence and its applications to estimation, control and learning under uncertain, situated conditions.

# 6 Neuroscience-Inspired Control for Mobile Navigation

## Problem description

At BIASlab we study how artificial agents can learn to control themselves in a (simulated) environment. We take inspiration from neuroscience and develop theory and (software) tools that allow for efficient, real-time operation in complex environments. This approach is named Active Inference (AIF), and applications range from robotics and navigation to personalized wearable devices.

The Digital Twin (DT) Lab creates models of complex processes and devices for the purpose of interactive exploration, training and visualization. Applications range from smart buildings and airplanes to factory floors and logistical processes.

The focus of the present project is on mobile navigation. You will work towards a demonstrator that integrates an AIF agent with a DT environment for mobile navigation. In order to achieve this, you will need to investigate to following:

1. How can an AIF agent be embedded in a DT environment?
2. How can a DT environment be employed as part of the AIF agent's control algorithm?

## Student task description

1. *Literature study on AIF*
   AIF is a relatively young but promising approach to control of artificial agents. The AIF approach is fundamentally probabilistic in nature. You will need to become familiar with Bayesian approaches to probabilistic inference and message passing algorithms on graphical models. It helps tremendously if you have a background in Bayesian machine learning (5SSD0 or similar).

2. *Become familiar with RxInfer*
   RxInfer is a package for efficient, reactive message passing on probabilistic models, written in the Julia programming language. This probabilistic programming package is developed in-house at BIASlab and enables real-time inference for control of AIF agents. Experience with proper (artificial intelligence) software development practice (5ARA0 or similar) is advised.

3. *Become familiar with DTs*
   You will need to become familiar with the mobile navigation DT environment. You will learn how to interact with this environment and understand its strengths and limitations.

4. *Build a prototype agent*
   In a first design iteration, the prototype agent will internalize a rudimentary model of the DT environment. However, inference on even a simple model might already lead to goal-directed navigation.

5. *Improve the agent*
   The prototype agent may then be improved by upgrading its probabilistic model. Specifically, you will investigate how (subroutines of) the DT environment itself can be used as part of the probabilistic model.

6. *Iterate* 4 and 5 until satisfaction and/or until time is up.

# 7 Football player simulation for coach assistance

## Problem description

This project is part of a long-term research program that aims to develop an intelligent agent (software) that provides strategic football (soccer) coaching decisions in real-time. We work together with the KNVB (Royal Dutch Football Association), and we aim to use the agent to assist coaching staffs of Dutch football teams.

As an example, consider an agent that receives the real-time coordinates of the players and the ball while the match is ongoing. The agent should be capable of analyzing the match in real-time and make coaching recommendations, e.g., switch our left-wing attacker to the right wing because the pairing with the opposing team's left-back is more advantageous.

Figure 3: Arjan Robben's big chance at the World Championship in 2010.

Intelligent decision making derives from our capacity to compare (the performance of) simulations of future scenarios. In this project, we aim to develop an agent that supports this type of "human intelligence"-inspired decision making. Technically, in this project you will explore the Free Energy Principle (FEP) as a guide on how to develop this type of agents. The FEP is a celebrated neuroscientific theory of how brains perceive and make decisions, based on a single objective, namely "Free Energy" minimization. As such, in a broader context, this project aims to develop a novel neuro-inspired approach to real-time intelligent decision making.

Clearly, developing such an agent is too difficult to accomplish in a single project. Therefore, this project is part of a larger program that encompasses multiple projects, where each project tries to accomplish intermediate milestones.

## Student task description

Your main job will be to develop an agent (a software program) that is able to cast strategic coaching decisions during a football game.

**Concrete tasks:**

- Specify the scope of the project. A first task is to specify what ”strategic coaching decisions” means in this project. We will need to describe and unpack this problem into a concrete and limited problem statement. This stage will be executed through discussions with your coaches and the KNVB.
- Do a literature study on existing approaches to data-driven coaching agents. Do a focused literature review and a select (at least 1) data-driven coaching algorithm from the literature that you will use as your ”reference” algorithm. Implement this reference algorithm.
- Do a literature study on the Free Energy Principle. In this project, our aim is to develop a new approach to data-driven coaching that learns through Bayesian machine learning methods. This can be described as an application of the Free Energy Principle (FEP). You will need to read about the FEP in general and about some previous efforts by BIASlab members and others on how to build FEP-inspired synthetic agents. Together with the literature study on existing coaching agents, this should lead to a small review paper (which we later embed in your thesis).
- Develop a FEP-based agent that learns to cast coaching decisions. The task of designing a new agent based on Bayesian machine learning may seem daunting, but you will be able to design your first agents after a few discussions with your supervisor.
- Evaluate your FEP-based agent in comparison to the reference method(s).
- Iterate steps 4 and 5 until satisfaction and/or time runs out.
- Write and defend an awesome thesis on a new data-driven approach to developing strategic football coaching agents.

## Project positioning

The Bayesian Intelligent Autonomous Systems lab (BIASlab) is a subgroup of the Signal Processing Systems (SPS) group that develops agents that interact with their environment through their sensors and actuators. The goal is to learn purposeful behavior, e.g., to navigate, play soccer or they may learn to decode speech signals under bad acoustic conditions. Our research projects are inspired by the latest insights from probabilistic machine learning, computational neuroscience and signal processing.

The KNVB is the Royal Dutch Football Association. You will be working with data scientists from the KNVB who have a tremendous amount of experience analyzing football matches and the behaviour of players. They have collected rich data sets on all matches played by the Dutch national team over the past few years. Their insights will be invaluable for the development of a useful football simulation model.

**Supervision:**

- Weekly progress meetings with dr. Wouter Kouw and prof. Wim Nuijten.
- Periodic progress meetings with contacts at the KNVB.
- The student should prepare update meetings, preferably with derivations or visualizations in interactive notebooks.

## Timeline

The project is available from September 2021 onwards. The total duration will be 32 weeks. Halfway through, there will be a "midterm" evaluation where the student must report on their activities and indicate how they will proceed for the remainder of the project. At the end, the student will write a thesis summarizing their work, their findings and possible future steps. The thesis will be presented in an official "defense" ceremony and a committee of experts will grade the student's work.

# 8 Active listening

## Problem description

The below excerpt is adopted from (Friston, 2021) and very nicely describes both the problem that we are tackling and the solution strategy:

*Speech recognition is not a simple problem. The auditory system receives a continuous acoustic signal and, in order to understand the words that are spoken, must parse a continuous signal into discrete words. To a naïve listener, the acoustic signal provides few cues to indicate where words begin and end (Altenberg, 2005; Thiessen and Erickson, 2013). Furthermore, even when word boundaries are made clear, there exists a many-to-many mapping between lexical content and the acoustic signal. This is because speech is not 'invariant' (Liberman et al., 1967)—words are never heard out of a particular context. An example is the identical perception of the phrases "grade A" and "grey day". When considering how words are generated, there is wide variability in the pronunciation of the same word among different speakers (Hillenbrand et al., 1995; Remez 2010)—and even when spoken by the same speaker, pronunciation depends on prosody (Bänziger and Scherer, 2005). From the perspective of recognition, two signals that are acoustically identical can be perceived as different words or phonemes by human listeners, depending on their context.*

In the seminal paper by Friston they describe a new approach and model to solve this speech recognition problem. They consider the task of setting the word boundaries and classifying the separated words as an active process, where they are performed simultaneously instead of sequentially. However, the model that they propose relies on an extensive set of (unnecessary) assumptions. In this project the goal is to reduce the set of assumptions, improve the model for acoustic sounds and to beat the performance of the model in (Friston, 2021). The baseline implementation has already been implemented by a previous master's student for a quick comparison.

## Student task description

The student will be working on improving the current implementation of the active listening paper. A thorough review of the used preprocessing steps and modeling choices needs to performed, such that it becomes clear where performance can be gained.

**Concrete tasks:**

- Familiarize yourself with the literature on probabilistic modeling and model comparison.
- Discuss problem and solution proposal with researchers at BIASlab.
- Improve the current model in terms of performance, scalability and preprocessing.

Figure 4: The ”carnavalsmachine” is a melodic speech synthesis device. It is capable of synthesising speech for composing Dutch music specially tailored for the annual religious festivities in the south of the Netherlands. Although its true operations are surrounded by secrecy, we have reasons to believe that a generative model is at the core of the machine. This generative model is capable of generating perceptually accurate speech and as the underlying model is probabilistic, it is also suited for recognizing speech through probabilistic inference. Researchers believe that the novel speech recognition approach of (Friston, 2021) provides key insights in the operations of this mysterious machine and that it helps improving its operations in the upcoming year. ©Lamme Frans

- Implement solution proposal using BIASlab resources.
- Experiment with the implementation and compare to current solutions.
- Analyze results and reflect on what has been achieved.
- Write a report detailing the advantages and limitations of this approach.
- Defend your work.

## Project positioning

The Bayesian Intelligent Autonomous Systems lab (BIASlab) is a subgroup of the Signal Processing Systems (SPS) group that develops agents that interact with their environment through their sensors and actuators. The goal is to learn purposeful behavior, e.g., to navigate, play soccer or they may learn to decode speech signals under bad acoustic conditions. Our research projects are inspired by the latest insights from probabilistic machine learning, computational neuroscience and signal processing.

**Supervision:**

- Weekly progress meetings with Bart van Erp.
- Weekly seminars with the BIASlab group, where we highlight our state-of-the-art research.
- The student should prepare update meetings, preferably with derivations or visualizations in interactive notebooks.
- All developed code should be accessible (e.g., on BIASlab's GitHub organization).

## Timeline

The project is available from March 11, 2022 onwards. The total duration will be 32 weeks. Halfway through, there will be a "midterm" evaluation where the student must report on their activities and indicate how they will proceed for the remainder of the project. At the end, the student will write a thesis summarizing their work, their findings and possible future steps. The thesis will be presented in an official "defense" ceremony and a committee of experts will grade the student's work.

## References

Friston, K. J., Sajid, N., Quiroga-Martinez, D. R., Parr, T., Price, C. J., Holmes, E. (2021). Active listening. Hearing Research, 399(Stimulus-specific adaptation, MMN and predicting coding), 107998. https://doi.org/10.1016/j.heares.2020.107998

# 9 Drift control in TwinScan lithography

## Problem description

TwinScan lithography systems are used at a very high speed to produce integrated circuits which requires sub-nanometer accuracy and precision. Due to hardware imperfections, these TwinScan systems drift, which causes a negative impact on both the (sub-)nanometer scale positioning and on the system throughput, as it requires (semi-manual) maintenance actions to reduce the impact of these drift.

The current drift controller (see Figure 5) is a static event-based feedback controller which gets activated only when it is necessary by means of a human request or when the system's internal maintenance kicks in. This leads to several challenging issues, e.g.,

1. The measurements are not ideal, i.e., they are subject to measurement noise which are of time-varying and colored nature, unpredictable jumps, and time lags that are introduced while performing measurements.

2. Drifts often possess a different dynamical structure over time for each machine. As a result, using a fixed-model structure is not optimal.

In this project you are challenged to design an improved drift controller based on modern machine learning technology.


Figure 5: TWINSCAN NXE:3350B is one of ASML's nanolithography systems.

## Student task description

The student will be challenged to use modern machine learning technology to develop a drift controller. The controller should learn online ("on-the-fly") from past measurements (while continuously providing feedback corrections) so as to cope with the time-varying dynamical nature of the drifts.

We expect that the proposed drift controller has several intrinsic characteristics:

- It should deal with non-uniform measurements in order to update the predictor's (internal) states.
- It should be rather general, in the sense that in principle any type of drift can be handled.
- It should possess a smart decision-making capability by learning from past measurements in order to differentiate between jumps, outliers and handle time-varying noises.
- It should not depend on a fixed-model structure (per class of drifts) as this would violate the previous requirement.

The performance of the proposed controller should be compared to the existing (non-adaptive) event-based drift controller, with both benchmark and real measurement data.

## Project positioning

The Bayesian Intelligent Autonomous Systems lab (BIASlab) is a subgroup of the Signal Processing Systems (SPS) group that develops agents that interact with their environment through their sensors and actuators. The goal is to learn purposeful behavior, e.g., to navigate, play soccer or they may learn to decode speech signals under bad acoustic conditions. Our research projects are inspired by the latest insights from probabilistic machine learning, computational neuroscience and signal processing.

ASML is a market and innovation leader in the semiconductor industry. They provide chipmakers with everything they need – hardware, software and services – to mass produce patterns on silicon through lithography. Their machines are incredibly precise, even up to nanometer-scale. ASML employs a large amount of industrial researchers and is quick to incorporate new technologies into their fabrication process.

## Supervision:

- Weekly progress meetings with Bert de Vries.
- Periodic progress meetings with Raaja Ganapathy Subramanian at ASML.
- The student should prepare update meetings, preferably with derivations or visualizations in interactive notebooks.

## Timeline

The project is available from Augustus 2020 onwards. The total duration will be 32 weeks. Halfway through, there will be a "midterm" evaluation where the student must report on their activities and indicate how they will proceed for the remainder of the project. At the end, the student will write a thesis summarizing their work, their findings and possible future steps. The thesis will be presented in an official "defense" ceremony and a committee of experts will grade the student's work.

# 10 Monitoring health through wearables

## Problem description

This is a project in collaboration with Philips on inferring a patient's health condition from wearable sensors. In the TRICA study, patients were sent home after elective surgery with two wearable devices, a patch and a wristband - see pictures, that measured several biometric signals (including heart rate, respiration rate, activity level (acticounts) and activity type (walking, running, cycling, other)). One goal of the TRICA study is to develop algorithms that, based on the measured data streams, detect in real-time when possible post-operative complications surface. This type of analysis is hindered by the fact that the recorded data is often unreliable or missing due to movement and other artifacts.


Figure 6: The TRICA study (which stands for Transitional Care) collects data from wearables for post-operative monitoring of recovery and potential complications.

In this graduation project you are challenged to develop a robust algorithm that, based on the recorded TRICA data base, forecasts heart rate and respiratory rate from activity level and type. In particular, we are interested in personalized algorithms that improve as more data is collected. In other words, your challenge is to develop an algorithm that both executes (predicts HR and RR) and learns to improve its performance as time moves on. The predicted HR and RR can then be used for forecasting health deterioration.

## Student task description

We expect the following tasks in this project:

- Specify the scope of the project.
  A first task is to specify the scope of the project. Together with you we

will develop the above description into a concrete problem statement that allows you to do a focussed literature review.

- Do a literature study on available methods.
  Do a focussed literature review and a select (at least 1) prediction algorithm from the literature that you will use as your "reference" algorithm. Implement this reference algorithm.

- Do a literature study on Bayesian machine learning-based prediction.
  At BIASlab, we develop Bayesian prediction algorithms such as advanced Kalman filters that map very nicely to the tasks in this project. You will study the literature in this field and summarize the literature and your plan in a small review paper (which you can later embed in your thesis).

- Develop an online Bayesian machine learning-based prediction algorithm for your task.

- Evaluate your algorithm in comparison to the reference method(s).

- Iterate steps 4 and 5 until satisfaction and/or time runs out.

## Project positioning

The project will be carried out at Philips Research and at BIASlab (FLUX-7.060), which is a subgroup of the Signal Processing Systems (SPS) group. BIASlab's research projects are inspired by the latest insights from machine learning, computational neuroscience and signal processing. You will get support (supervision) both from a senior Philips researcher and BIASlab researcher.

### Supervision

- Weekly progress meetings with Bert de Vries.

- Periodic progress meetings with Reinder Haakma at Philips.

- The student should prepare update meetings, preferably with derivations or visualizations in interactive notebooks.

## Timeline

The project is available from October 2020 onwards. The total duration will be 32 weeks. Halfway through, there will be a "midterm" evaluation where the student must report on their activities and indicate how they will proceed for the remainder of the project. At the end, the student will write a thesis summarizing their work, their findings and possible future steps. The thesis will be presented in an official "defense" ceremony and a committee of experts will grade the student's work.

# 11 Graphical editor for building agents

## Problem description

You are challenged to develop a graphical web-based editor/debugger for developing and monitoring the dynamics of simulated agents. Please have a look at this video by Bret Victor that demonstrates the essence of the envisioned editor. Rather than developing electronic circuits, in this project we focus on developing intelligent agents that minimize a quantity known as Free Energy through message passing in a factor graph. See the 5SSD0 Bayesian Machine Learning & Information Processing lesson on intelligent agents for more information.

In the second part of the thesis work, you use the editor/debugger to analyze the internal dynamics (i.e., monitor the free energy distribution) of a few well-known (but not well-understood) intelligent agents, thus leading to an increased understanding of how Free Energy minimizing dynamics give rise to intelligent behavior (see e.g. this video for a neuroscience perspective.)

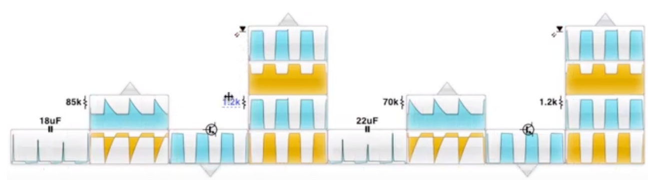


Figure 7: Visualization of (electronic circuit) dynamics (source.

## Student task description

Please note that, aside from having a demonstrated interest in dynamic systems modeling and/or machine learning, for this project it is essential that you also have good programming skills, in particular related to web front-end technologies. The project is ideally suited to a good programmer with interests in visual editing and machine learning. It is also possible to do this project with 2 people.

The task is organized into the following steps:

- Scope: a first task is to narrow the scope of this project. Development of a good editor/debugger may take years, so we will focus on a few simple features that showcase proof-of-concept and allow dynamic monitoring of free energy distribution during simulations.
- Literature study of modern supporting technology for development of web-based editors.
- Literature study of free energy minimization of dynamic systems in factor graphs, finalized by an interim report and project planning.

- Develop/update a simple web-based graphical editor/debugger based on Bret Victor's principles.
- Analyze a highly performant but not well-understood intelligent agents by online monitoring and manipulation of the internal dynamics of these agents.
- Iterate steps 4 and 5 until satisfaction and/or time runs out.
- Write and defend an awesome thesis on a new approach to interpretable AI through online visualization and manipulation of the simulated system.

## Project positioning

The Bayesian Intelligent Autonomous Systems lab (BIASlab) is a subgroup of the Signal Processing Systems (SPS) group that develops agents that interact with their environment through their sensors and actuators. The goal is to learn purposeful behavior, e.g., to navigate, play soccer or they may learn to decode speech signals under bad acoustic conditions. Our research projects are inspired by the latest insights from probabilistic machine learning, computational neuroscience and signal processing.

### Supervision:

- Weekly progress meetings with Bert de Vries.
- The student should prepare update meetings, preferably with derivations or visualizations in interactive notebooks.
- All developed code should be accessible (e.g., on BIASlab's Github organization).

## Timing

The project is available from April 26th 2020 onwards. The total duration will be 32 weeks. Halfway through, there will be a "midterm" evaluation where the student must report on their activities and indicate how they will proceed for the remainder of the project. At the end, the student will write a thesis summarizing their work, their findings and possible future steps. The thesis will be presented in an official "defense" ceremony and a committee of experts will grade the student's work.

# 12 Reinforcement learning for home security systems at Signify

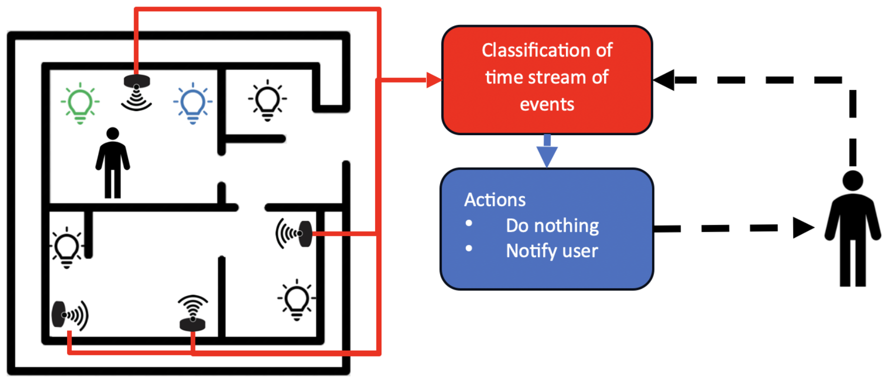


## General Description

Signify-Philips Hue has launched a home security proposition. Motion sensors, cameras, and contact sensors can be used to detect a possible intruder. These sensors send discrete events to a centralized logic (bridge) where a decision on a possible intrusion needs to be made. A requirement with a home security proposition is to keep the rate for false positives at a minimum. This is challenging as some sensors may trigger on unintended signals such as a PIR sensor triggering on clouds moving in front of sun, motion triggers from moving plants due to wind, and false object detection by cameras. In this internship the candidate is expected to develop a reinforcement learning probabilistic model that based on sensor events makes decision on possible intrusions. The method should learn by itself when and how much each sensor can be trusted. Possible actions the logic can take are generating notifications to the user and generating a full alarm. An extra ingredient that may be studied is to include a human happiness model. A human may be happy to see that the alarm gets triggered upon entering an armed area. A human may be unhappy to get a full alarm when nothing is going on. The choice for actions of a system may be optimized for human happiness. Yet another ingredient is to explore active inference options. Instead of optimizing for immediate human happiness, one may optimize actions for a future happiness of the user.

## Task description

- Familiarize yourself with Bayesian probabilistic models
- Get a dataset, either from a toy-model or from the company Signify

- Develop an algorithm that adapts to the reliability of sensors and makes decisions on possible intrusions while respecting requirements of false and true positives
- Implement algorithm in BIASlab and demonstrate improvements
- Extend approach by including a human happiness model and optimize actions for immediate human happiness (Bonus)
- Extend approach by optimizing actions for future human happiness (Bonus)