WHITEPAPER

Transforming healthcare with Reinforcement Learning

August 2022

OVERVIEW

Table of contents

PEOPLE

White paper contributors

ALI EL HASSOUNI
PhD candidate @ VU Amsterdam,
(former) Data Scientist, Mobiquity

ADI GASKELL
Forbes Contributor

CHAPTER 1

Foreword

Thank you for your interest in our “Transforming healthcare with Reinforcement Learning” white paper.

In healthcare, patients are interested in understanding their own health to improve their well being, and businesses need to respond to these new expectations with an optimised standard of patient care.

Following their experience with digital consumer offerings, patients expect a personalised approach across all areas of their life. Through digital technology, we can achieve personalisation in healthcare, creating personalised, human centered health experiences that empower people to take greater control of their health and wellbeing.

One of the most promising technologies to achieve this is Artificial Intelligence (AI), in particular Reinforcement Learning (RL), which is used in certain consumer industries, for example to improve Facebook in-app recommendations for users or optimising advertising on Microsoft.

By adopting Reinforcement Learning, MedTech and pharma companies can optimise healthcare processes for the long term, using data to improve patient outcomes as digital health solutions continue to evolve.

Our white paper provides you with a comprehensive view of how RL can be used to provide value to your business, along with further reading for those interested in the nuts and bolts of how it works.

We would urge you to get in touch if you would like to explore how this approach can help your business in its digital journey.

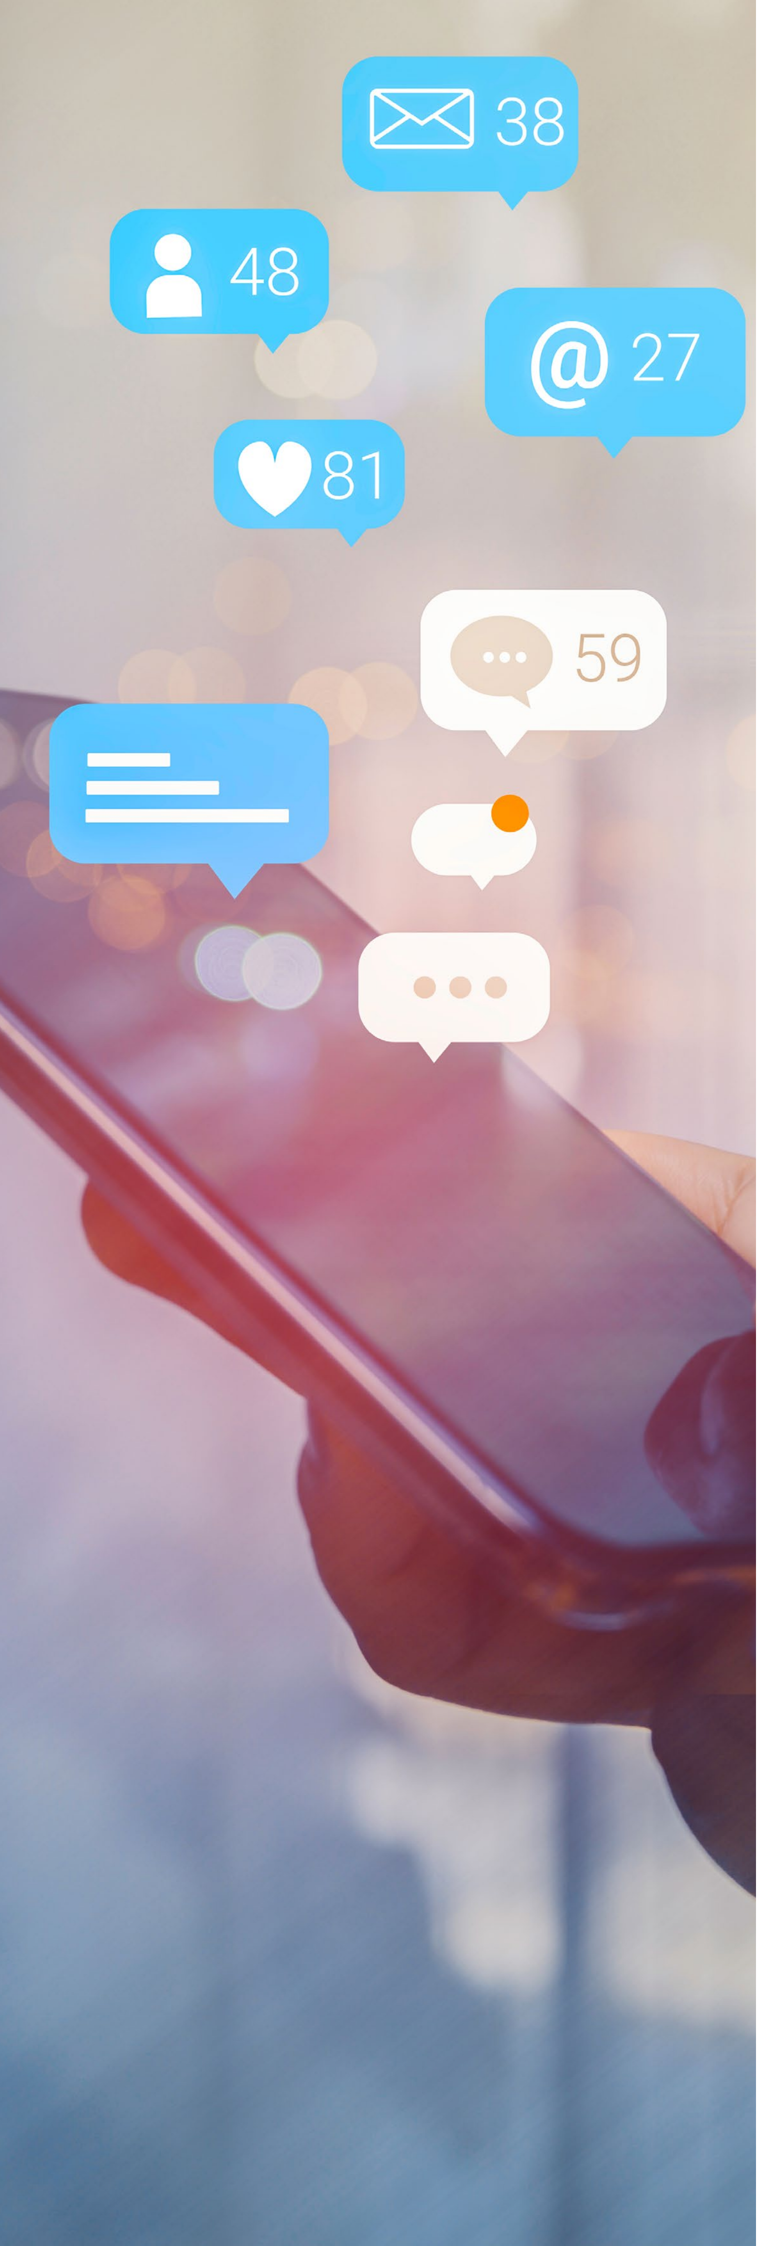

CHAPTER 2

Why is Reinforcement Learning useful for MedTech and pharma companies?

Reinforcement Learning is an advanced type of Machine Learning where a computer is taught by experience to act optimally in an environment, constantly improving processes to achieve a long-term goal. Put simply, the learning process finds and learns a sequence of actions to deliver a desired outcome for the end-user. This holds great value for health-tech and MedTech businesses to navigate multi-stakeholder journeys and improve patient outcomes.

To date, Reinforcement Learning with Model Personalisation is applied mostly in a consumer context where there is lots of data, and in a risk free experimental environment with clear business objectives on what to optimise. Practical use cases have focused on online marketing, advertising or optimising notification messages, like Facebook, which has used Reinforcement Learning to improve in-app recommendations - providing an ultra personalised experience for users.

To learn more about Reinforcement Learning with Model Personalisation please see the Technical Glossary below.

The most famous example of Reinforcement Learning is in gaming, with AlphaGo: A computer programme trained to play the Chinese board game 'Go', considered the most complex and challenging strategy game, with more possible board configurations than the number of atoms in the known universe. It was considered impossible for a computer to beat a human in Go due to the game's complexity, until AlphaGo was developed and became the first computer to defeat a world champion human in Go in 2015.

THE OPPORTUNITY FOR REINFORCEMENT LEARNING IN HEALTHCARE

In healthcare, there is an opportunity to scale Reinforcement Learning, as there is a clear objective in healthcare that needs to be optimised constantly: Improving patient outcomes. Across the multi-stakeholder journey, companies are striving towards this objective, but there is not always a straightforward approach to achieve it. That's where Reinforcement Learning can help.

For example, chronic disease management and clinical care are in essence a sequence of clinical interventions over time aiming to cure or increase a patient's quality of life. A Reinforcement Learning model is able to consider the long-term effects of treatment to an individual, using the patient's state to inform care and treatment plans. While a doctor bases this on their experience, a Reinforcement Learning model is capable of determining time-dependent decisions for the best treatment for a patient.

The use of Reinforcement Learning in healthcare also enables improvement of long-term outcomes by factoring in any potential delayed effects of treatments.

For businesses, applying Reinforcement Learning to digital health solutions has clear benefits, driving value directly to the patient and personalising healthcare through Machine Learning methods. As digital healthcare continues to evolve post-COVID-19, adopting sophisticated methods such as Reinforcement Learning in healthcare will be a gamechanger for the industry. RL is enabling pharma and MedTech companies to design and deliver scalable digital health solutions that improve patient outcomes, reducing costs by integrating clinical workflows - increasing the amount of time healthcare professionals can spend with their patients.

THE FUTURE OF HEALTHCARE: CLINICAL APPLICATIONS OF REINFORCEMENT LEARNING

There is a golden opportunity for pharma and MedTech companies to push the boundaries of digital health through emerging methodologies in the healthcare domain such as Reinforcement Learning.

Machine Learning as is, is being used in the industry, however, the adoption of Reinforcement Learning remains low.

Academic examples of how Reinforcement Learning could be used in healthcare are becoming popular with data scientists and AI researchers, providing an opportunity for businesses to open up revenue streams while delivering better patient outcomes.

SEPSIS

Researchers at Imperial College London developed an Artificial Intelligence system that could help treat patients with sepsis - a common and, if not treated appropriately, deadly infection leading to acute organ dysfunction.

A WHO report on the burden of sepsis reports 49 million individuals affected worldwide, of which 20 million are children younger than 5. In the UK, sepsis kills around 44,000 people every year.

In addition, the financial impact of sepsis is significant, with approximately 1.7 million individuals affected in the USA, costing $24 billion in total annual hospital costs, at a cost of $22k per patient.

The system learnt the best treatment strategy for a patient by analysing the records of 100,000 US hospital patients in intensive care units. It examined every single doctor's decision affecting the patients, comparing the AI recommendations with the standard care and how that would have impacted patient outcomes.

The management of intravenous fluids and vasopressors in sepsis is a key clinical challenge. Besides general guidelines, no tool exists to personalise treatment of sepsis and assist clinicians in making decisions in real time and at the patient level that could improve this situation. This leads to clinical variability in sepsis treatment and probably therefore suboptimal decisions leading to poorer outcomes for the patients.

As a result, the Reinforcement Learning model provides individual and clinically interpretable treatment decisions for sepsis that could improve patient outcomes if scaled.

MOODBUSTER

Moodbuster is another example of how data-driven platforms can improve patient outcomes using Reinforcement Learning. A European project and not-for-profit tool, Moodbuster is used to treat psychological problems remotely.

One solution to scale the Moodbuster application was through Reinforcement Learning, creating an AI therapist to nudge users in an optimised way to improve their physical and mental health. Using an e-health simulator with VU Amsterdam, researchers fine tuned an algorithm to build a messaging service as part of Moodbuster to support people with mental health problems. This project is currently at a clinical trial stage with patients in the Netherlands.

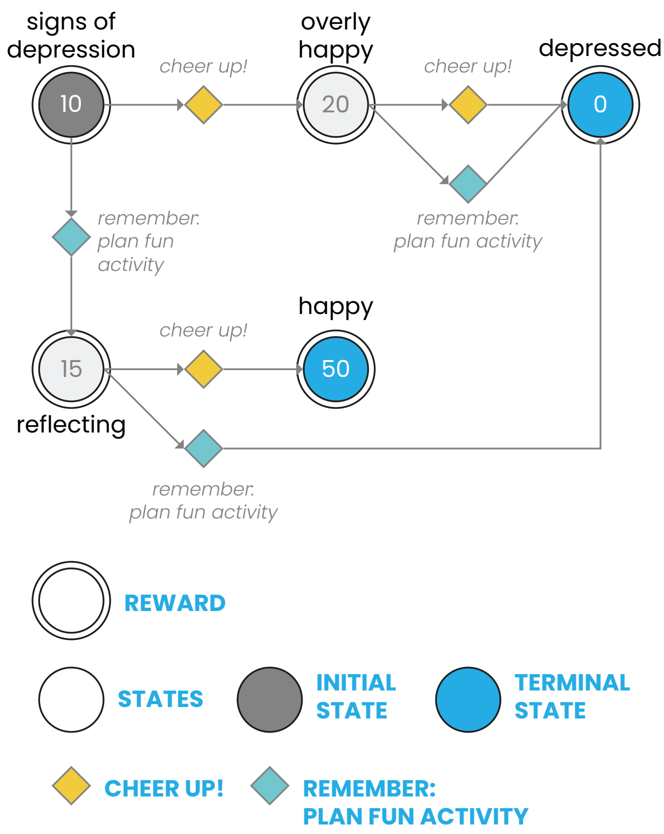

CHAPTER 3

How to harness Reinforcement Learning in healthcare

For MedTech and pharma companies to harness the benefits of Reinforcement Learning, we recommend the following steps:

1. START SIMPLE

From a practical point of view, RL can be challenging to integrate immediately. Start with a smaller amount of training data to help drive adoption.

2. BUILD FOR THE FUTURE

As Reinforcement Learning models continuously learn from their predictions and new data, a good scalable base infrastructure is essential. This will provide an opportunity to learn what kind of data the model will need to improve, building its capability for the future.

3. COLLABORATE WITH HEALTHCARE SPECIALISTS

As the RL model will require clinical data to learn and evolve, integrating healthcare expertise is essential, and best achieved by a close collaboration between data scientists and healthcare professionals.

4. DATA IS KING

The right type of data is key for developing the right policies. Once your Reinforcement Learning model has matured, leveraging the right data is needed to achieve a personalised approach in order to deliver value in healthcare.

5. COLLABORATE WITH USERS

Even the best advice will not have any impact if the recipient is not able or willing to listen. So, make sure that the end-user, i.e. patients or healthcare professionals, are included in every step of the development process: From problem definition, to solution design, to testing and adoption in a clinical setting.

CHAPTER 4

Our approach to Reinforcement Learning

Our approach to applying Reinforcement Learning in healthcare drives value directly for healthcare providers and MedTech companies by improving patient outcomes, applying a personalised, proactive data analysis model using Reinforcement Learning to predict the sequence of actions needed to achieve the desired health outcomes.

We drive value by developing analytics that can underpin mobile applications and strive to improve the level of personalisation by means of novel machine learning methods that:

- Intelligently select, from a predefined set, interventions to support a particular user.
- Do so proactively, e.g. on the basis of predictions of user status.

An app would learn when and how to apply which interventions that best suit a particular patient, for a specific goal or purpose, and it would learn as it obtains more experiences and interactions from the patient.

This should lead to:

- Better personalisation for the patient
- More engaged patients
- Improved patient satisfaction and well being
- New forms of revenue generation for MedTech and pharmaceutical companies

CHAPTER 5

Current and future applications of Reinforcement Learning

Reinforcement Learning has been applied in some clinical areas, such as in telehealth, remote medicine, clinical medicine, personalised healthcare, and across a range of mHealth disciplines in areas such as fitness and coaching apps.

While the potential uses are indeed varied, the main applications fall into three main categories: dynamic treatment regimes in chronic disease or critical care, automated medical diagnosis, and other general domains such as health resources allocation and scheduling, optimal process control, drug discovery and development, as well as health management.

The early use cases that are emerging in the field also highlight the challenges in deploying Reinforcement Learning successfully. Many of the early breakthroughs in Reinforcement Learning have leaned heavily on Deep Learning.

As a subset of Machine Learning, Deep Learning is where artificial networks use algorithms inspired by the human brain, to learn from large amounts of data unsupervised. As it requires a huge amount of data with which to train the algorithm, deep Reinforcement Learning also faces similar challenges while also requiring the computational power to crunch through a huge number of possible iterations.

One way around this is to utilise more traditional algorithms to underpin the Reinforcement Learning feedback loop. These kinds of models are able to generalise effectively without requiring huge quantities of data to operate with. To unlock the potential of this approach fully, we can start using Deep Reinforcement Learning to replace more traditional approaches once enough data is available.

Healthcare also presents evident ethical challenges in terms of training the model as it is clearly not possible to train the system and to perform exploration on actual patients, so a realistic simulation is required to allow the model to run through numerous iterations. It's for this reason that video games have proven a popular testing ground for many Reinforcement Learning applications to date, as they provide a ready made virtual environment for the models to play around in – incorporating a 'gamification' model to keep the user engaged. In healthcare, we can train on offline data, using off-policy Reinforcement Learning to use transfer learning and further train with online data.

Just as digital twins are growing in popularity in the manufacturing industry, it might be that digital twin simulators may be required for Reinforcement Learning to truly take off in healthcare as the experiments could then be performed at scale without affecting any real patients.

CONTACT US

Using Reinforcement Learning to solve business frictions

Interested in learning how RL in healthcare can help you with creating new revenue streams, remove the frictions in the multi-stakeholder journey, and deliver a personalised patient experience?

Contact us at: marketing.eu@mobiquityinc.com

TECHNICAL GLOSSARY

More about Reinforcement Learning

In discussions around Artificial Intelligence, a lot of confusion can arise due to the seemingly interchangeable use of various terms. While it's easy to conflate Machine Learning, Deep Learning, and Reinforcement Learning they are in fact all slightly different approaches to Artificial Intelligence.

The first distinction is around how the machine learns. For much of the history of Artificial Intelligence rule-based programming was used to allow machines to draw inferences from data. Machine Learning differs from this in that it allows the machine to learn on its own from the data it is fed.

Deep Learning is commonly thought of as a subset of Machine Learning with the main training algorithms being neural networks. In practise these algorithms often require large datasets to train good models. For example, training an image classification model to detect leukemia, a dataset of blood cell images that have been tagged by human doctors as either leukemic or non-leukemic is required.

Reinforcement Learning works best when you know what it is you want the system to achieve but are not so sure on the best way to achieve that outcome. It requires you to be able to quantify the variables in the environment and define the "reward function" of the system so it can successfully learn from each iteration. As such, you need to be able to have full information about the environment the system will operate in, and be able to clearly define what a "good" outcome will look like.

PERSONALISATION

Personalisation in Reinforcement Learning is the process of developing policies that are tailored towards the individuals or patients with the goal of achieving the highest impact possible.

Most traditional Machine Learning models follow a one-size fits all approach using a generalizing algorithm. This often results in a shorter and faster learning period and a model that is good enough for most use cases, but this comes at a cost: lots of data. A more iterative learning process like with Reinforcement Learning requires less data to start with and when more and more data becomes available, the model will get better and better, outperforming the more general model in the long run. A grouped policy is the midway between a one-size fits all approach and separate models. Through personalising in clusters of users, a grouped policy means a higher personalisation level and faster convergence.

When combined with a more personalised approach (not a single model for all, but multiple models for clusters of patients, or ideally a unique model per user), Reinforcement Learning becomes more valuable over the long term.

POOLED

GROUPED

SEPARATE

A technical deep dive

RL UNIQUE CHARACTERISTICS

There are some unique characteristics that set Reinforcement Learning apart from other Machine Learning models. There is no supervisor (a “teacher” showing the model what good looks like), but just a delayed reward signal. It learns a policy (for instance a set of rules) which determines a sequence of actions. The actions of the agent influence the subsequent data it receives from the environment.

REWARDS

The goal of the agent is to perform sequences of actions to maximize the cumulative future reward. A reward is a feedback signal that indicates how well the agent is doing. Actions may have consequences that are only visible in the long run, leading to a delayed reward. The agent has to make a trade-off between immediate and long-term rewards.

REINFORCEMENT LEARNING

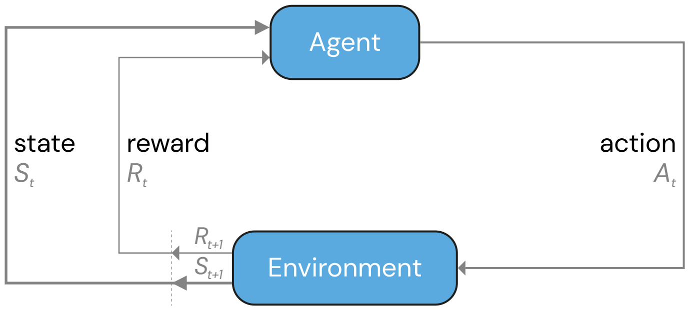

Agent	
Performs an action	A_t
Receives data about the environment state	S_t
Receives a reward	R_t

Environment	
Receives an action, performed by the agent	A_t
Shows state	S_{t+1}
Provides reward	R_{t+1}

This leads to a sequence of actions, state observations and rewards which we call history.

History = $S_1, R_1, A_1, \ldots, S_{t-1}, R_t, A_t$

LEARNING PROCESS

Reinforcement Learning is a machine learning process of developing a policy that learns to act optimally in a certain environment. Technically the policy is a mapping function from the state information to an action. The mapping can be deterministic action = Policy(state) or stochastic Policy(action|state).

One approach to RL is using value functions. The value function predicts the total future reward. The value function is used to select between actions by evaluating the expected total future reward given the current state. γ is the discount rate that tells how important the immediate and the long-term rewards are $V_{\pi}(s)=E[R_{t+1}+\gamma R_{t+2}+\gamma^2 R_{t+3}+\dots \mid S_t = s]$

A model predicts the output of the environment at the next time step. It predicts the state at t+1 step given the state and action at t and it predicts the reward at t+1 given the state and action at t.

LEARNED POLICY

The agent tries to learn an optimal policy, by maximizing the future reward using known information (exploitation) combined with its experience with the environment (exploration). It is important to find a good balance between exploration and exploitation. It is hard to address this in a pure mathematical manner. As such a RL model combined with a large dataset is a better approach.

FURTHER READING

Hoogendoorn, M., Funk, M. 2018. Machine Learning for the Quantified Self: https://www.springer.com/gp/book/9783319663074

Roggeveen, L. el Hassouni, A. Ahrendt, J. Guo, T. Fleuren, L. Thoral, P. Rj Girbes, A. Hoogendoorn, M. WG Elbers, P. 2021. Transatlantic transferability of a new reinforcement learning model for optimizing haemodynamic treatment for critically ill patients with sepsis: https://www.sciencedirect.com/science/article/pii/S0933365720312689

Komorowski, M., Celi, L.A., Badawi, O., Gordon, A.C. and Faisal, A.A., 2018. The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. Nature medicine, 24(11), pp.1716-1720. https://www.nature.com/articles/s41591-018-0213-5

Premier Inc. Analysis: Hospital-Associated Sepsis Decreased by 15% from 2015-2018 https://www.premierinc.com/newsroom/pressreleases/premier-inc-analysis-hospital-associated-sepsis-decreased-by-15-from-2015-2018

WHO. Global report on the epidemiology and burden of sepsis https://www.who.int/publications/i/item/9789240010789

Vincent, J.L., Marshall, J.C., Ñamendys-Silva, S.A., François, B., Martin-Loeches, I., Lipman, J., Reinhart, K., Antonelli, M., Pickkers, P., Njimi, H. and Jimenez, E., 2014. Assessment of the worldwide burden of critical illness: the intensive care over nations (ICON) audit. The lancet Respiratory medicine, 2(5), pp.380-386. https://www.ncbi.nlm.nih.gov/pubmed/24740011

el Hassouni, A., Hoogendoorn, M., Eiben, A.E. and Muhonen, V., 2020, July. Structural and Functional Representativity of GANs for Data Generation in Sequential Decision Making. In International Conference on Machine Learning, Optimization, and Data Science (pp. 458-471). Springer, Cham. https://link.springer.com/chapter/10.1007/978-3-030-64583-0_41

DE GEUS, S.V., EL HASSOUNI, A., DOBLER, D., MUHONEN, V., TIMMERS, D. and PRICE, R., 2020. Utilizing Available Data to Warm Start Online Reinforcement Learning. https://science.vu.nl/en/Images/stageverslag-geus_tcm296-946174.pdf

el Hassouni, A., Hoogendoorn, M., van Otterlo, M., Eiben, A.E., Muhonen, V. and Barbaro, E., 2018. A clustering-based reinforcement learning approach for tailored personalization of e-Health interventions. arXiv e-prints, pp.arXiv-1804. https://arxiv.org/ abs/1804.03592

den Hengst, F., Grua, E.M., el Hassouni, A. and Hoogendoorn, M., 2020. Reinforcement learning for personalization: A systematic literature review. Data Science, (Preprint), pp.1-41. https://content.iospress.com/articles/data-science/ds200028

El Hassouni, A., Hoogendoorn, M., Eiben, A.E., Van Otterlo, M. and Muhonen, V., 2019, October. End-to-end Personalization of Digital Health Interventions using Raw Sensor Data with Deep Reinforcement Learning: A comparative study in digital health interventions for behavior change. In 2019 IEEE/WIC/ACM International Conference on Web Intelligence (WI) (pp. 258-264). IEEE. https://ieeexplore.ieee.org/abstract/document/8909419

el Hassouni, A., Hoogendoorn, M. and Muhonen, V., 2018, October. Using generative adversarial networks to develop a realistic human behavior simulator. In International Conference on Principles and Practice of Multi-Agent Systems (pp. 476-483). Springer, Cham. https://link.springer.com/chapter/10.1007/978-3-030-03098-8_32

Saria, S. Individualized sepsis treatment using reinforcement learning. Nat Med 24, 1641–1642 (2018). https://doi.org/10.1038/s41591- 018-0253-x

Peine, A., Hallawa, A., Bickenbach, J. et al. Development and validation of a reinforcement learning algorithm to dynamically optimize mechanical ventilation in critical care. npj Digit. Med. 4, 32 (2021). https://doi.org/10.1038/s41746-021-00388-6

Soder, H.E., Cooper, J.A., Lopez-Gamundi, P. et al. Dose-response effects of d-amphetamine on effort-based decision-making and reinforcement learning. Neuropsychopharmacol. (2020). https://doi.org/10.1038/s41386-020-0779-8

Komorowski, M., Celi, L.A., Badawi, O. et al. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. Nat Med 24, 1716–1720 (2018). https://doi.org/10.1038/s41591-018-0213-5

CONTACT

Get in touch

Interested in learning how Reinforcement Learning in healthcare could help you with creating more personalised healthcare solutions? Please reach out for a call with us.

NIels Stol,

Healthcare Lead EMA, Mobiquity

+31 6 83 33 9349

nstol@mobiquityinc.com

/in/nielsstol

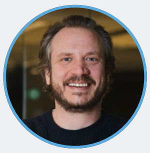

Teun Schutte,

Strategy director | Practice lead

Life science & Healthcare, Mobiquity

+31 6 82 75 7033

tschutte@mobiquityinc.com

/in/teunschutte