

VU Research Portal

A model a day keeps the doctor away

el Hassouni, Ali

2022

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

el Hassouni, A. (2022). *A model a day keeps the doctor away: Reinforcement learning for personalized healthcare*.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Contents

Acknowledgements	3
Summary	5
Samenvatting	7
1 Introduction	11
1.1 Motivation	12
1.2 Research Questions	15
1.3 Scope	15
1.4 Contributions	19
1.5 List of Papers	21
1.6 Extra Publications	24
I Reinforcement Learning for Personalization	25
2 A Systematic Literature Review	27
2.1 Introduction	28
2.2 Reinforcement Learning for Personalization	31
2.3 Algorithms	35
2.4 A Classification of Personalization Settings	47
2.5 A Systematic Literature Review	50
2.6 Results	55
2.7 Discussion	64
2.8 Acknowledgements	66
2.9 Appendix A. Queries	67
2.10 Tabular View of Data	69
3 Personalization: Pooled, Grouped or Separate?	79
3.1 Introduction	80
3.2 Related Work	83
3.3 Methodology	85
3.4 Evaluation Environments	91
3.5 Experimental Setup	99

3.6	Results	106
3.7	Discussion	114
3.8	Acknowledgment	115
II	Bridging the Reality Gap Safely	117
4	Bridging the Reality Gap with GANs	119
4.1	Introduction	120
4.2	Related Work	122
4.3	Model-based Reinforcement Learning	123
4.4	Simulation Environment	125
4.5	Using GANs to Generate Sensor Data	127
4.6	Experiments	128
4.7	Results	130
4.8	Conclusion	131
5	Structural and Functional Representativity	133
5.1	Introduction	134
5.2	Related Work	136
5.3	Method	137
5.4	Experimental Setup	142
5.5	Results	146
5.6	Conclusion	153
III	End-to-End Personalization	155
6	End-to-End Deep Reinforcement Learning	157
6.1	Introduction	158
6.2	Methods	160
6.3	Experimental Setup	165
6.4	Results	170
6.5	Conclusion	174
IV	Real-world Applications	175
7	Transferability of Reinforcement Learning Models for Sepsis	177
7.1	Introduction	179
7.2	Background and Related Work	180
7.3	Methods	191

7.4	Results	201
7.5	Discussion	214
7.6	Conclusion	217
7.7	Data Preprocessing	217
7.8	State Space Features	219
8	pH-RL: A Personalization Architecture to Bring Reinforcement Learning to Health Practice	221
8.1	Introduction	222
8.2	Related Work	224
8.3	pH-RL - A Framework for Personalisation with RL	227
8.4	Real-world Performance Evaluation	235
8.5	Results	237
8.6	Discussion	241
8.7	Acknowledgements	242
8.8	Appendix	242
V	Conclusion	245
9	Conclusion	247
9.1	Discussion	248
9.2	Research Questions	252
9.3	Future Work	253
	List of Figures	255
	List of Tables	261
	SIKS Dissertatiereeks	311