

Visual explanations for AI decisions

Fostering trust with transparency and control



Jeroen Ooge
jeroenooge.com

KU LEUVEN



Katrien Verbert
augment.cs.kuleuven.be

KU LEUVEN

Vraag: Wat is de hoofdstad van de staat Florida?

Antwoord: Tallahassee

Kwaliteit afleiders: Goed


Keuzelijst afleiders: Miami, New York, Los Angeles, Limburg, Noorwegen, Zuidpool

visualization

interaction

shepherding

direct explanation



gevorderde beginner | Volg mij is dit nu je level voor het onderwerp Hoofdbewerkingen

Welke moeilijkheidsgraad wil je voor de volgende oefeningenreeks?

Heel makkelijk | Makkelijk | Gewoon | Moeilijk | Heel moeilijk

ingen in de reeks juist oplost, dan stijgt je level:

Expert

Bedreven

Competent

Gevorderde beginner

Beginner

Start de reeks

Aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

- Oefening 37
- Oefening 26
- Oefening 21

Hoe is je nieuw niveau bepaald?

Wiski schat jouw niveau en de moeilijkheid van oefeningen in. Beide veranderen bij het je niveau is gestegen na het maken van de reeks oefeningen. Daarna is het nog extra gestegen door je feedback.

Maak oefening 21

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht



How good do you think you are at mathematics?

Expert: mathematics holds no secrets for you

Proficient: you score better than average on math

Competent: you score average on mathematics

Advanced beginner: basic exercises are no problem

Novice: you often have a hard time understanding

Submit


Patrol ID: 3033 | Age: 47 | Blood Sugar: 7.8 g/dl | Drinking Status: This Standard | Physical Activity Level: Low

Blood Sugar: 7.5 | Weight Measure: 112 | BMI: 33.9 | Physical Activity Level: Low




Recommendations to reduce risk: Exercise everyday for 30 min, Before walk seriously by 14 min

Impactful Risk Factors: Blood Sugar, Weight Measure, BMI, Physical Activity Level



This is an AI model



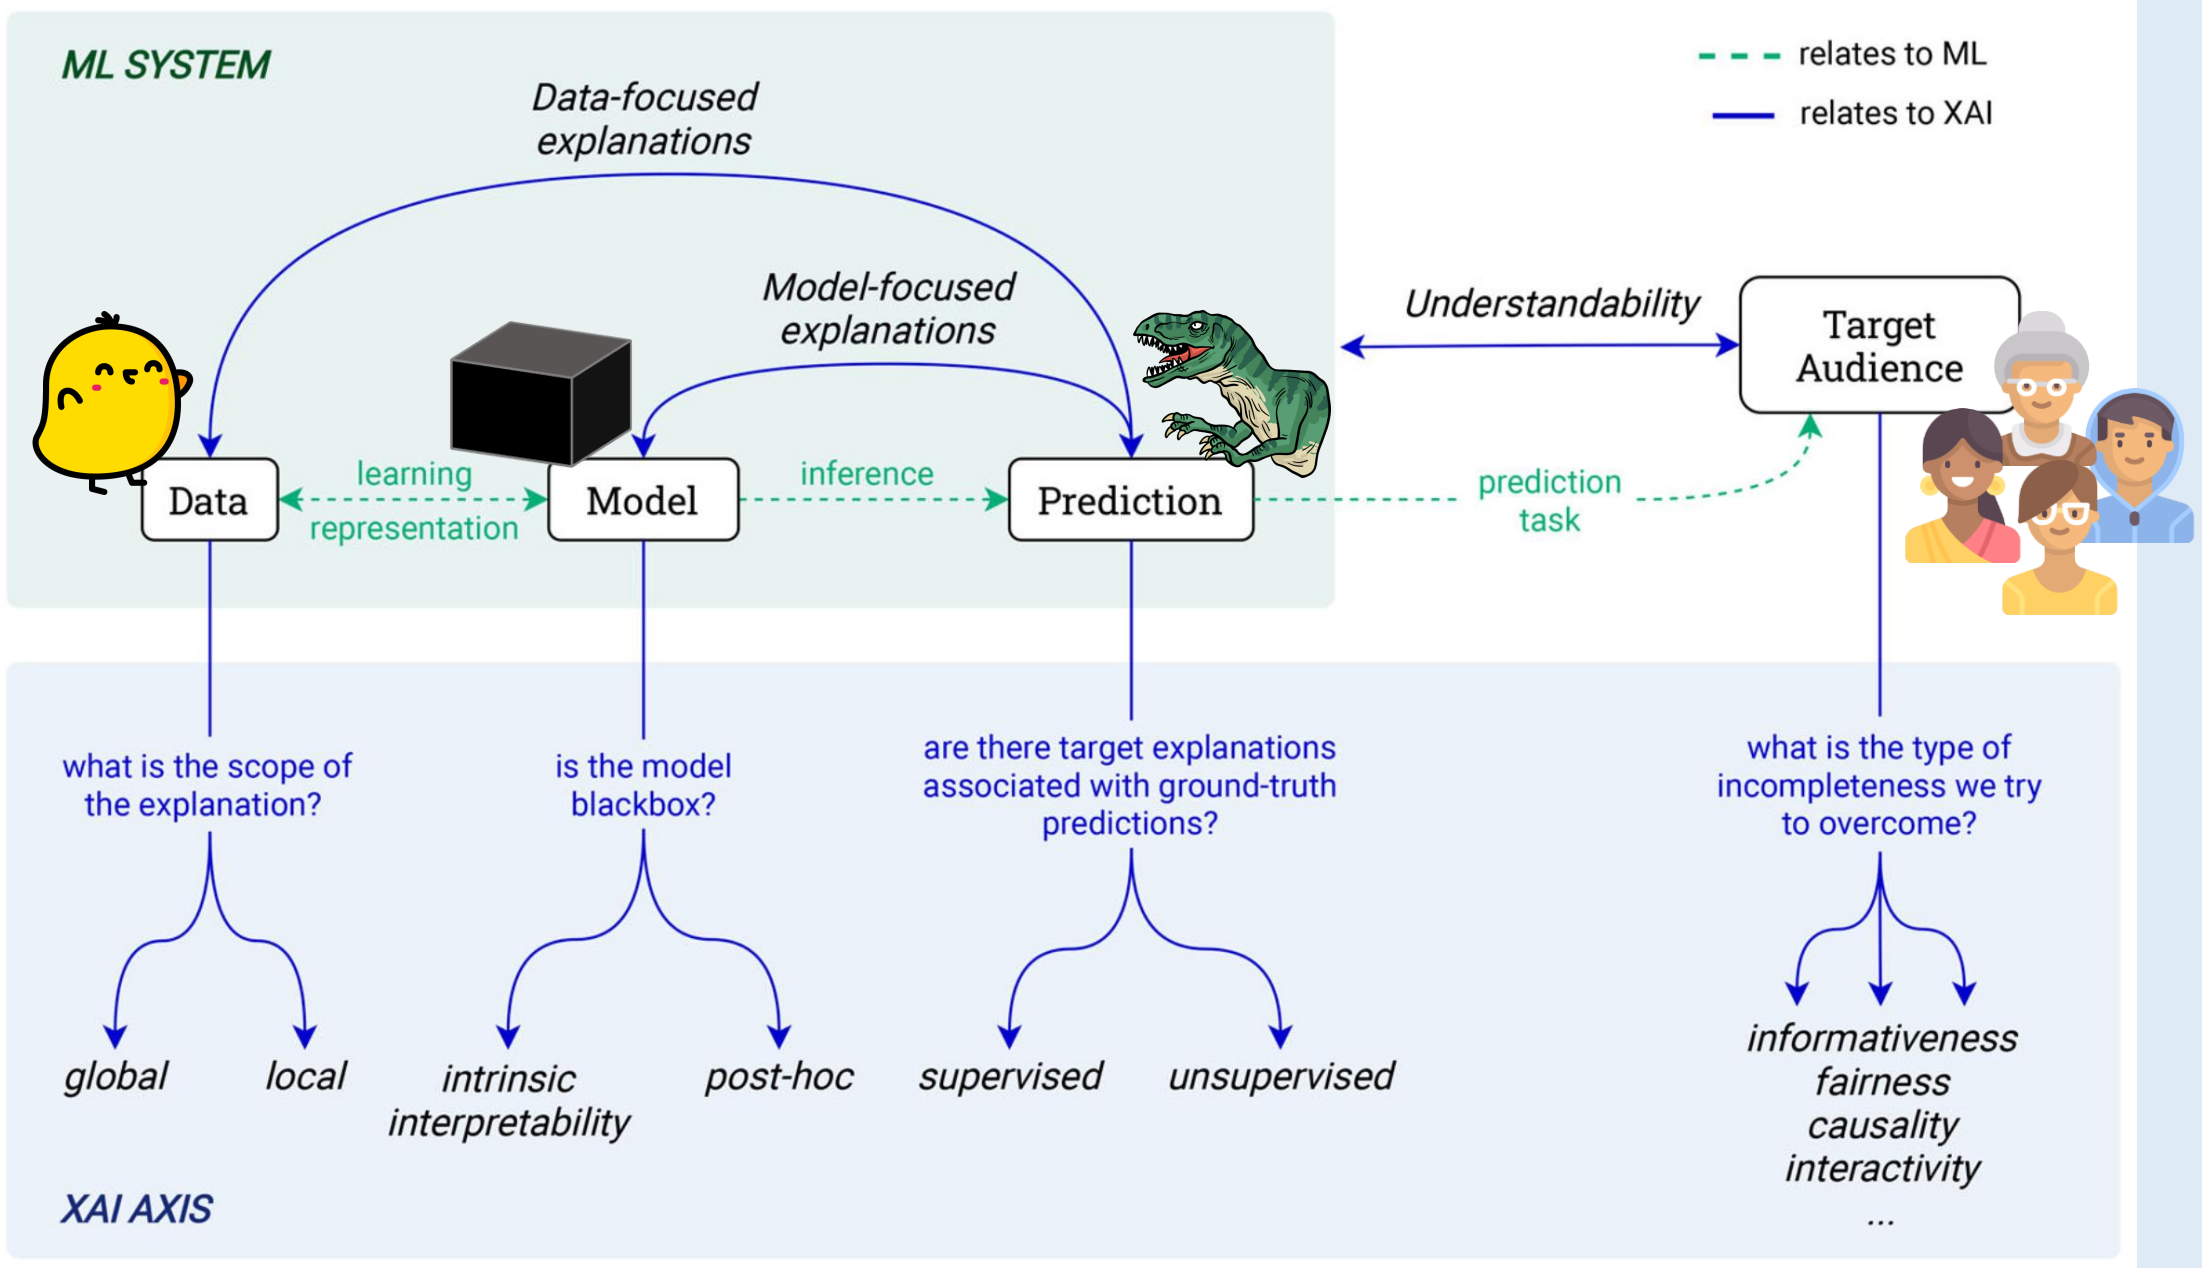
Input



AI model

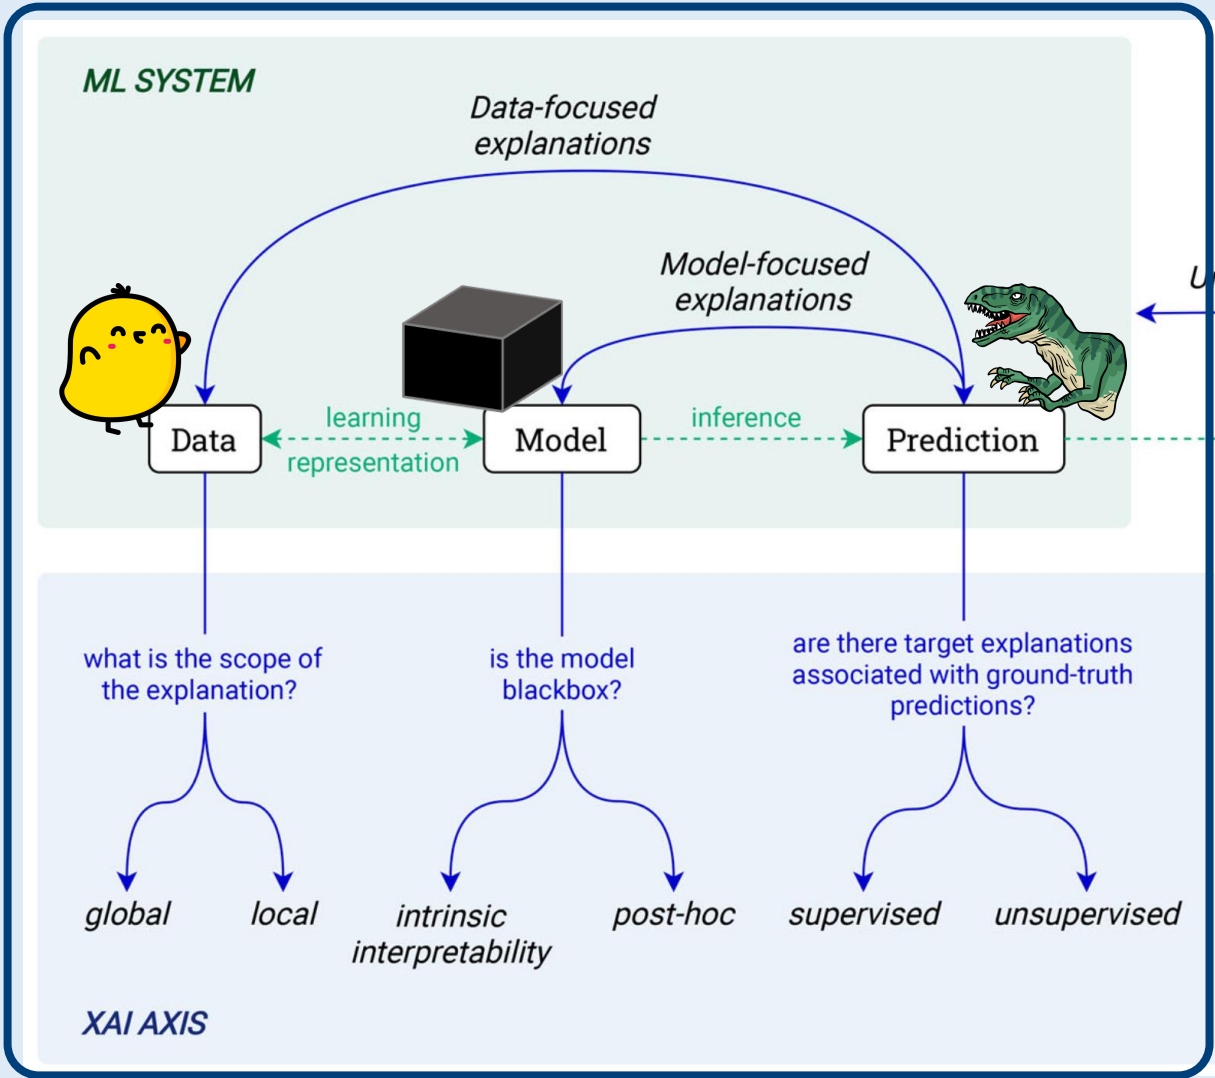


Output

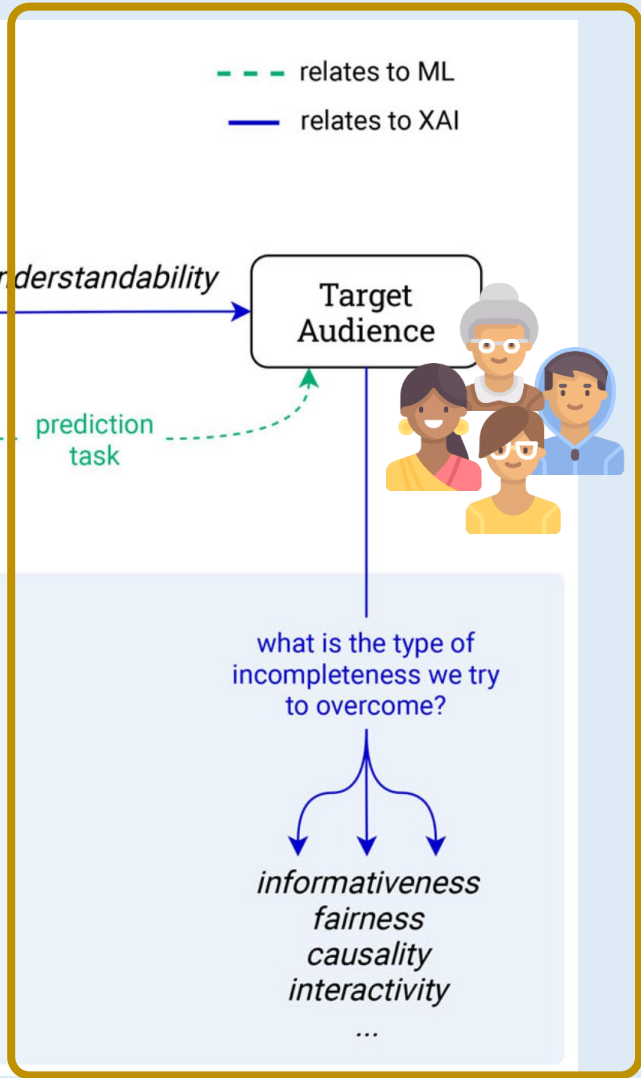


Darius Afchar, Alessandro Melchiorre, Markus Schedl, Romain Hennequin, Elena Epure, and Manuel Moussallam. 2022. Explainability in Music Recommender Systems. *AI Magazine* 43, 2: 190–208. <https://doi.org/10.1002/aaai.12056>

Algorithmic XAI approaches




Human-centred XAI approaches




Q. Vera Liao and Kush R. Varshney. 2022. Human-Centered Explainable AI (XAI): From Algorithms to User Experiences. <https://doi.org/10.48550/arXiv.2110.10790>

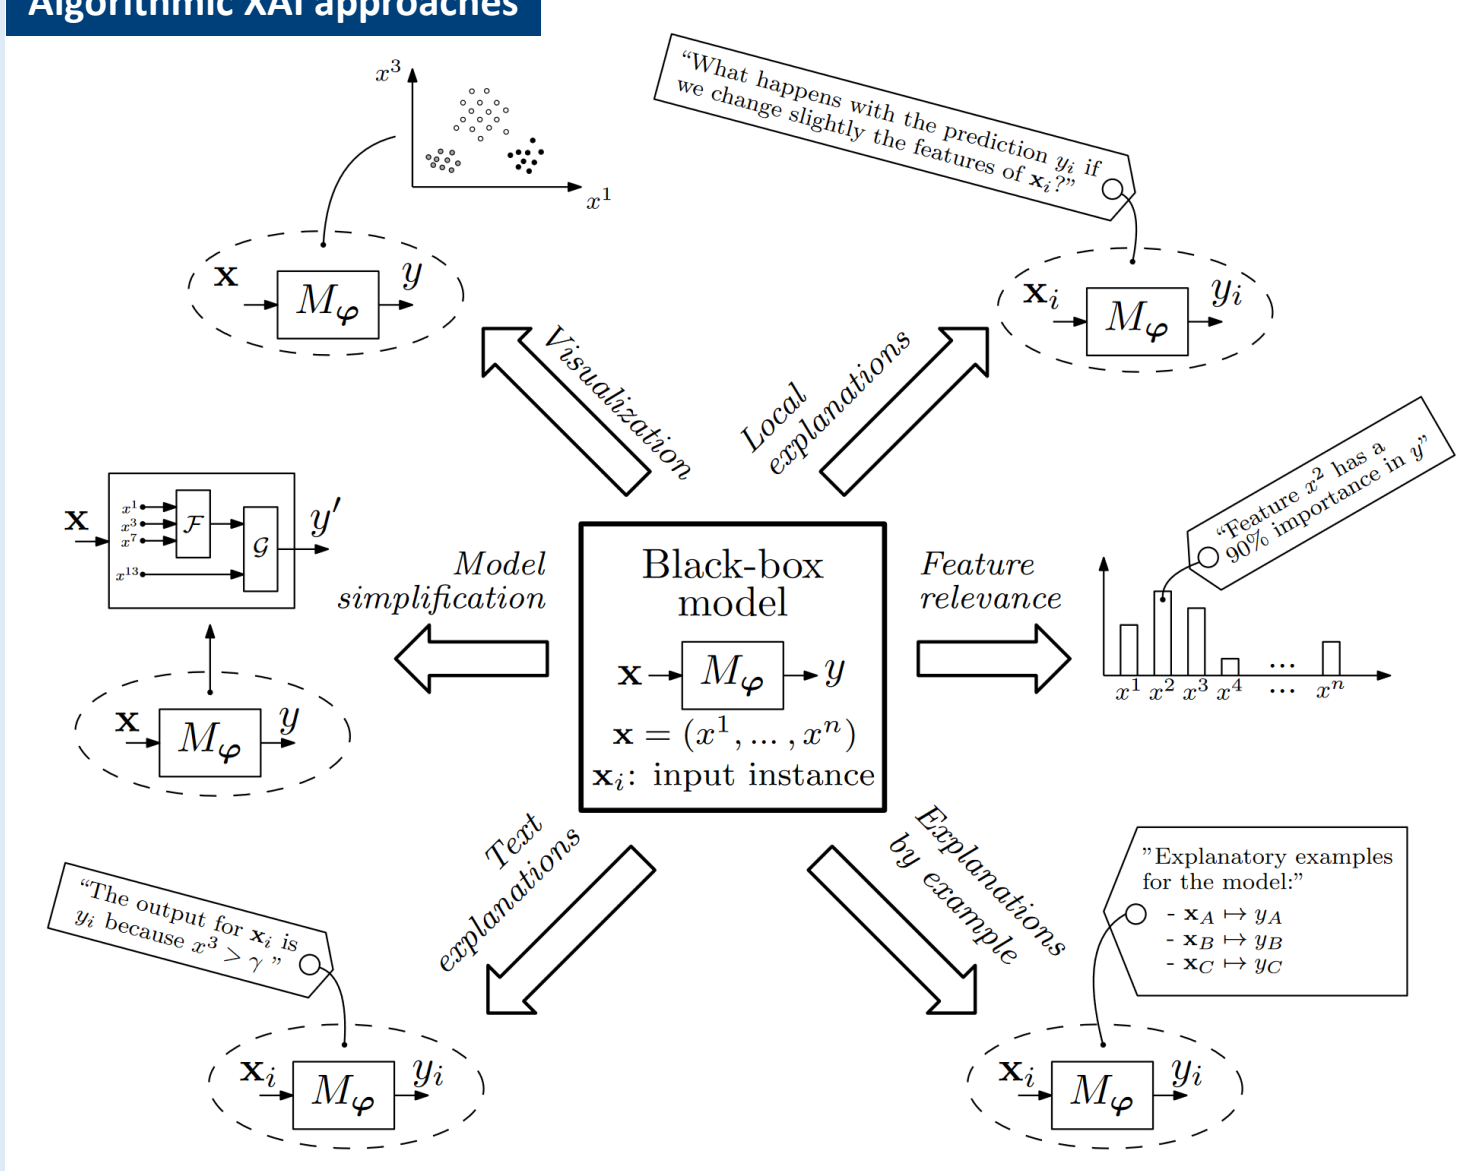
Algorithmic XAI approaches



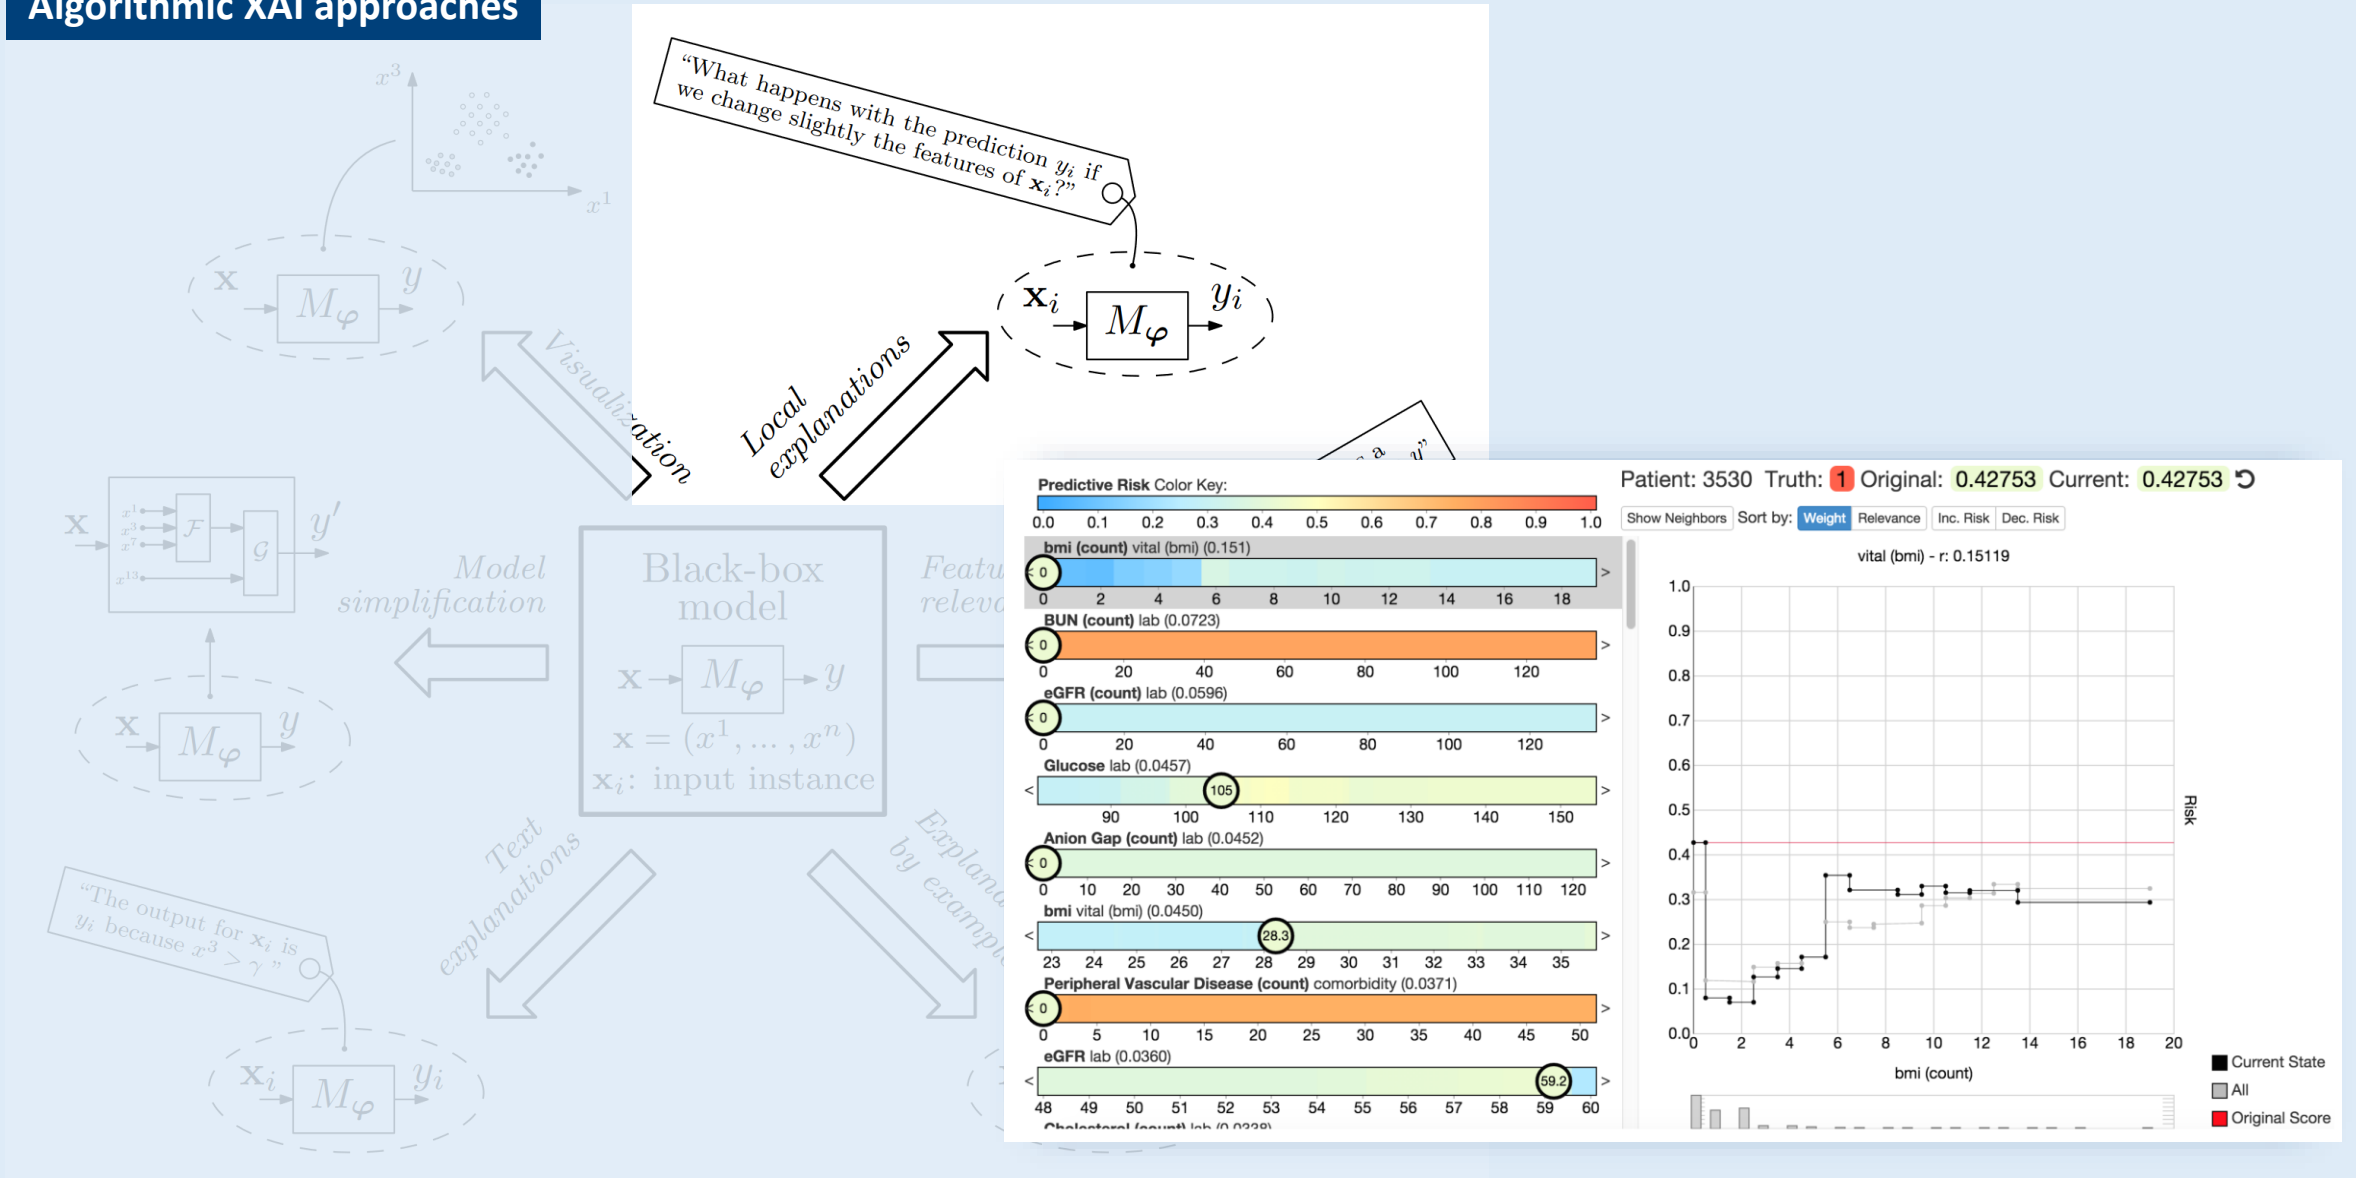
Human-centred XAI approaches



Algorithmic XAI approaches




Algorithmic XAI approaches



J. Krause, A. Perer, and K. Ng. 2016. Interacting with predictions: Visual inspection of black-box machine learning models. 5686–5697.


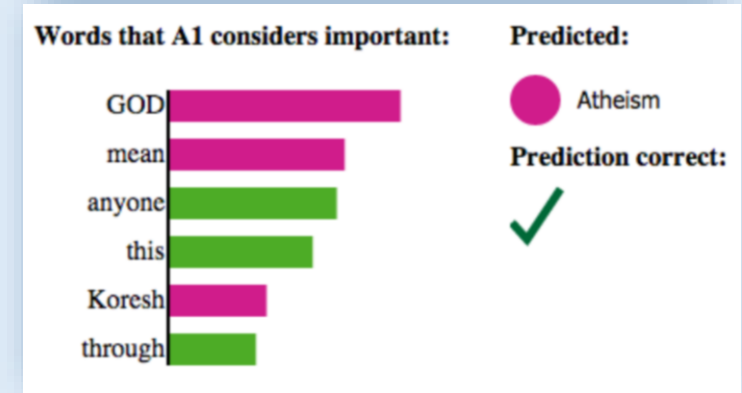
<https://doi.org/10.1145/2858036.2858529>

Algorithmic XAI approaches

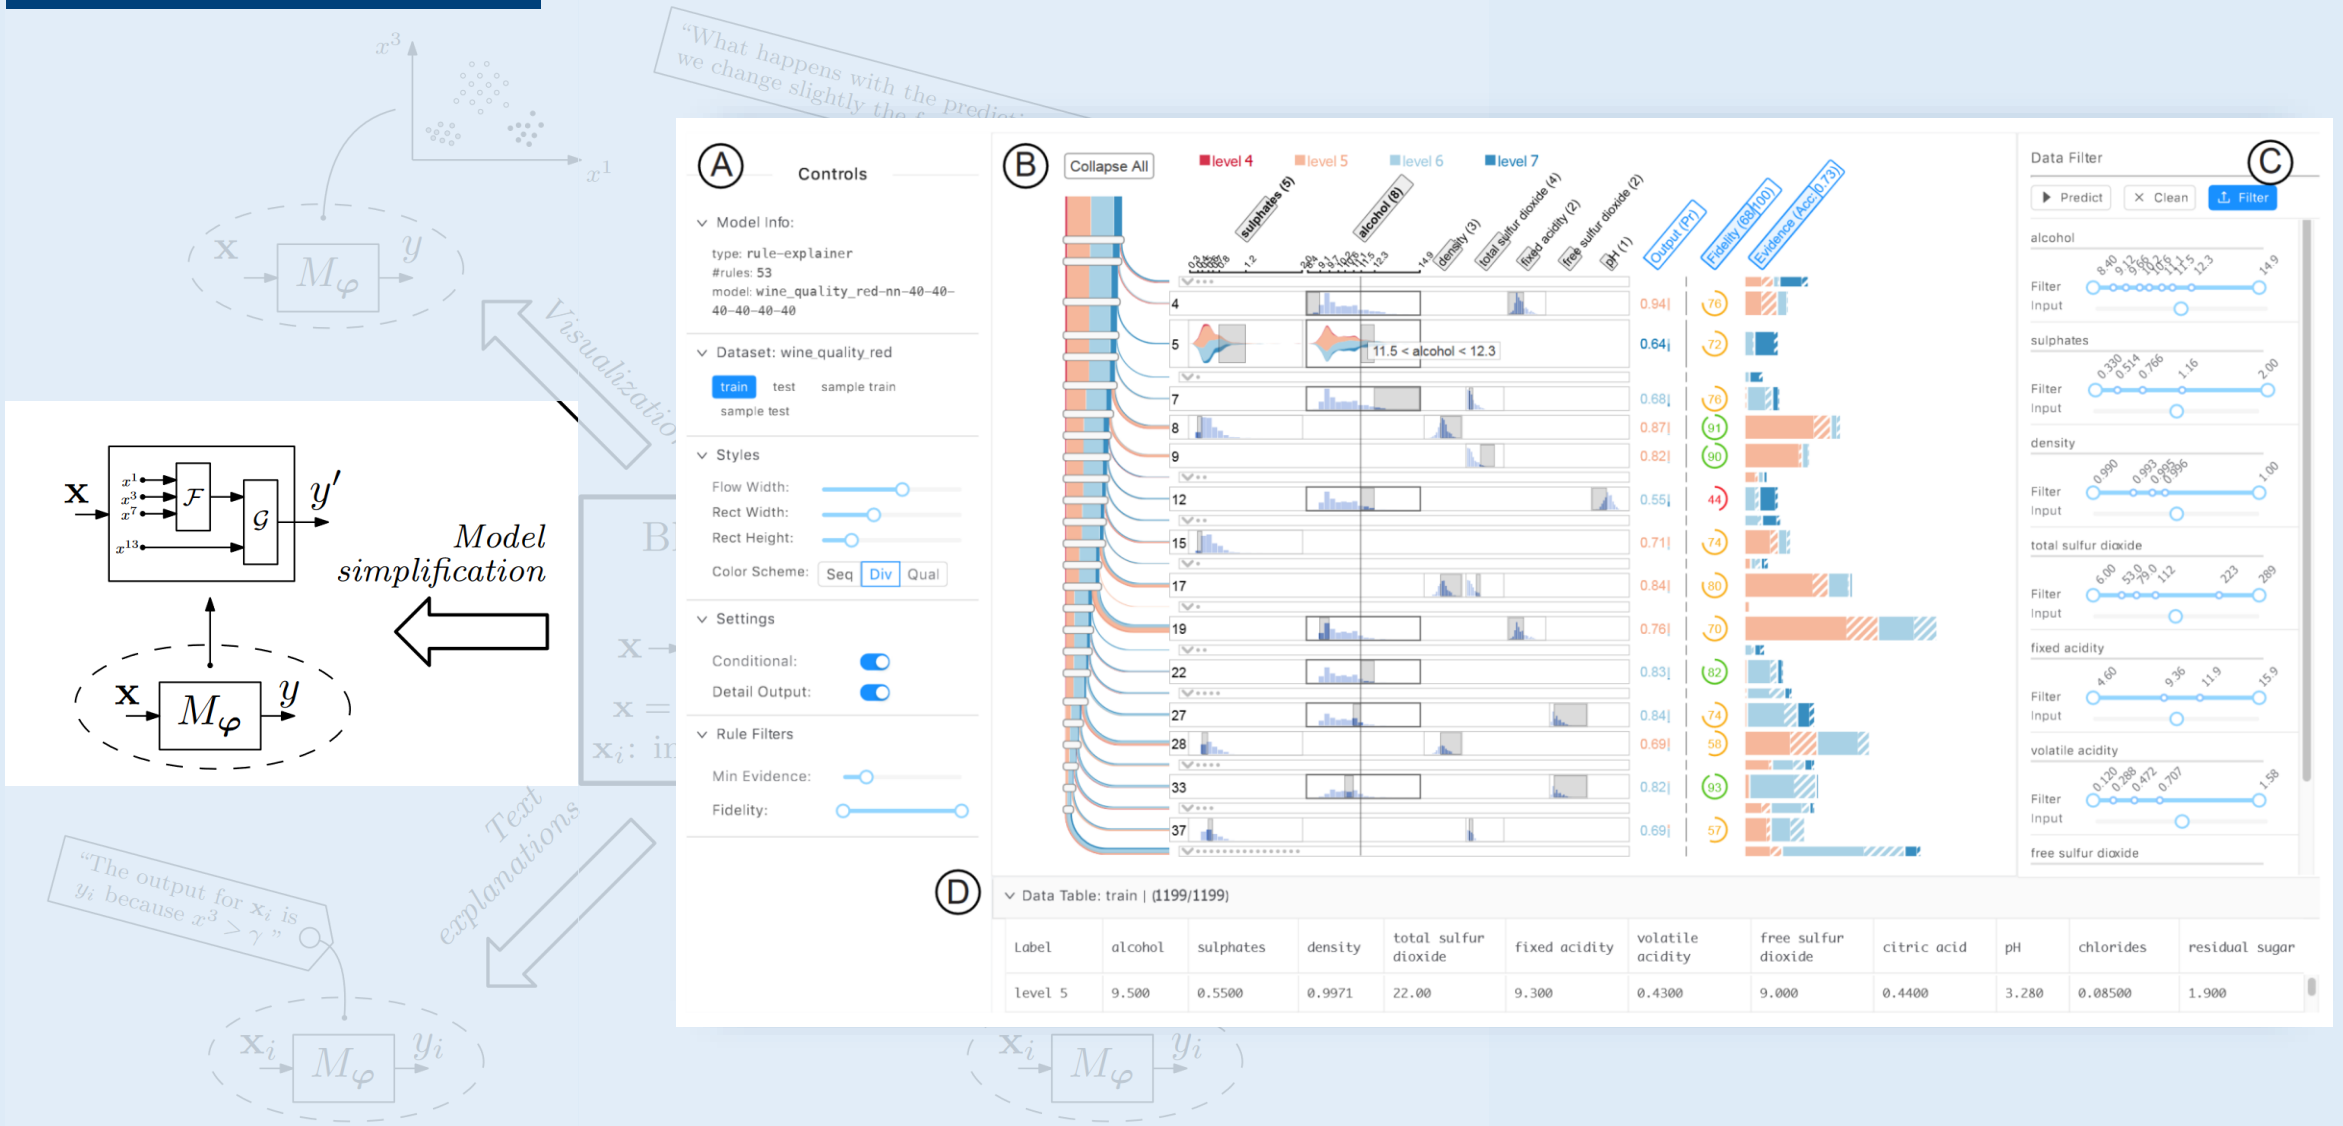


Scott M Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems*.
<https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf>

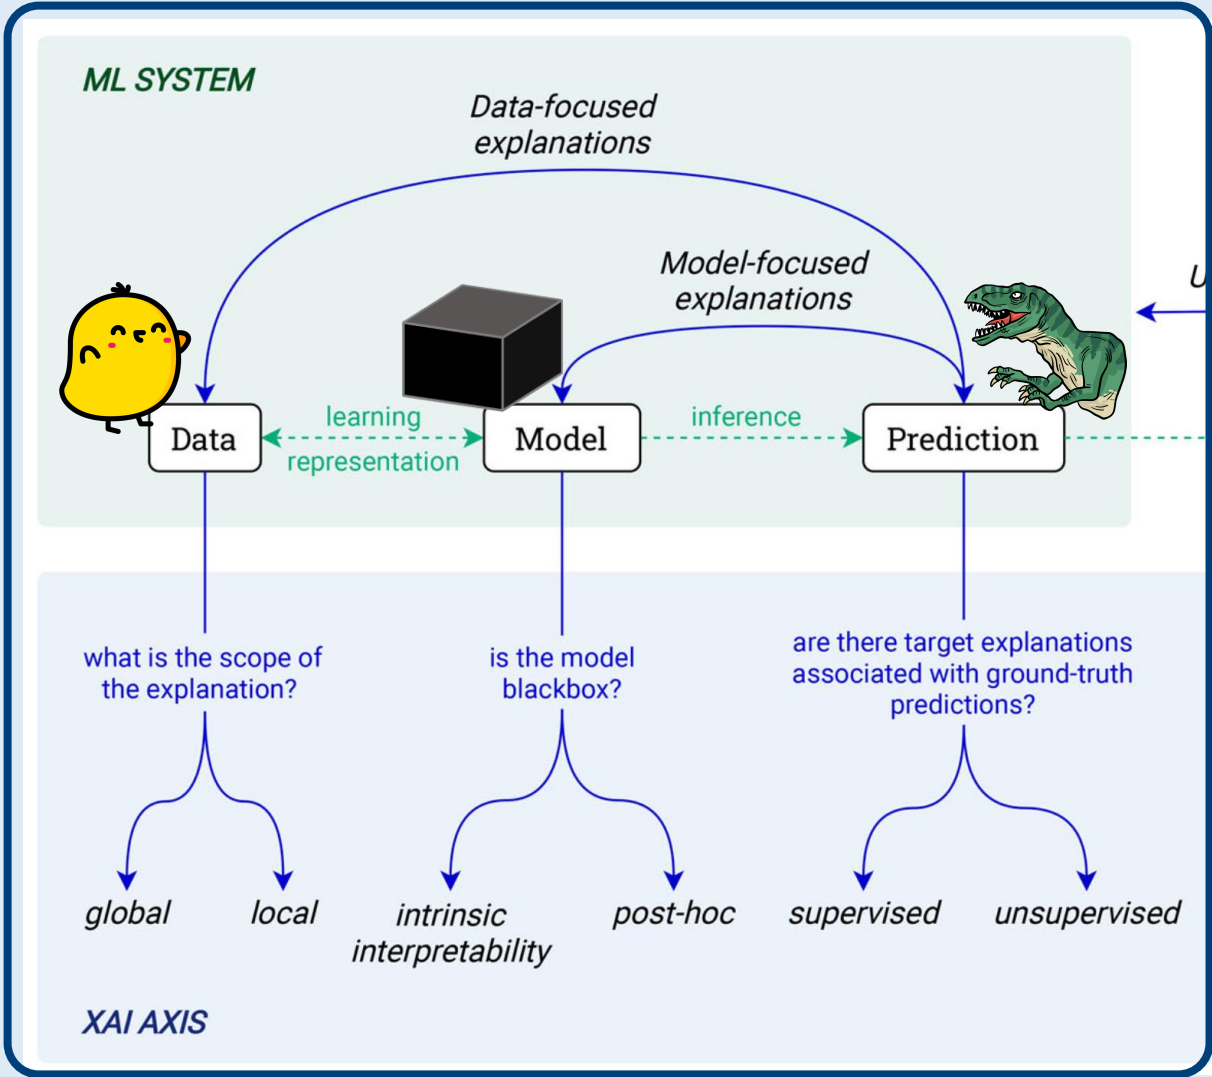
Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*, 1135–1144. <https://doi.org/10.1145/2939672.2939778>




Algorithmic XAI approaches



Algorithmic XAI approaches

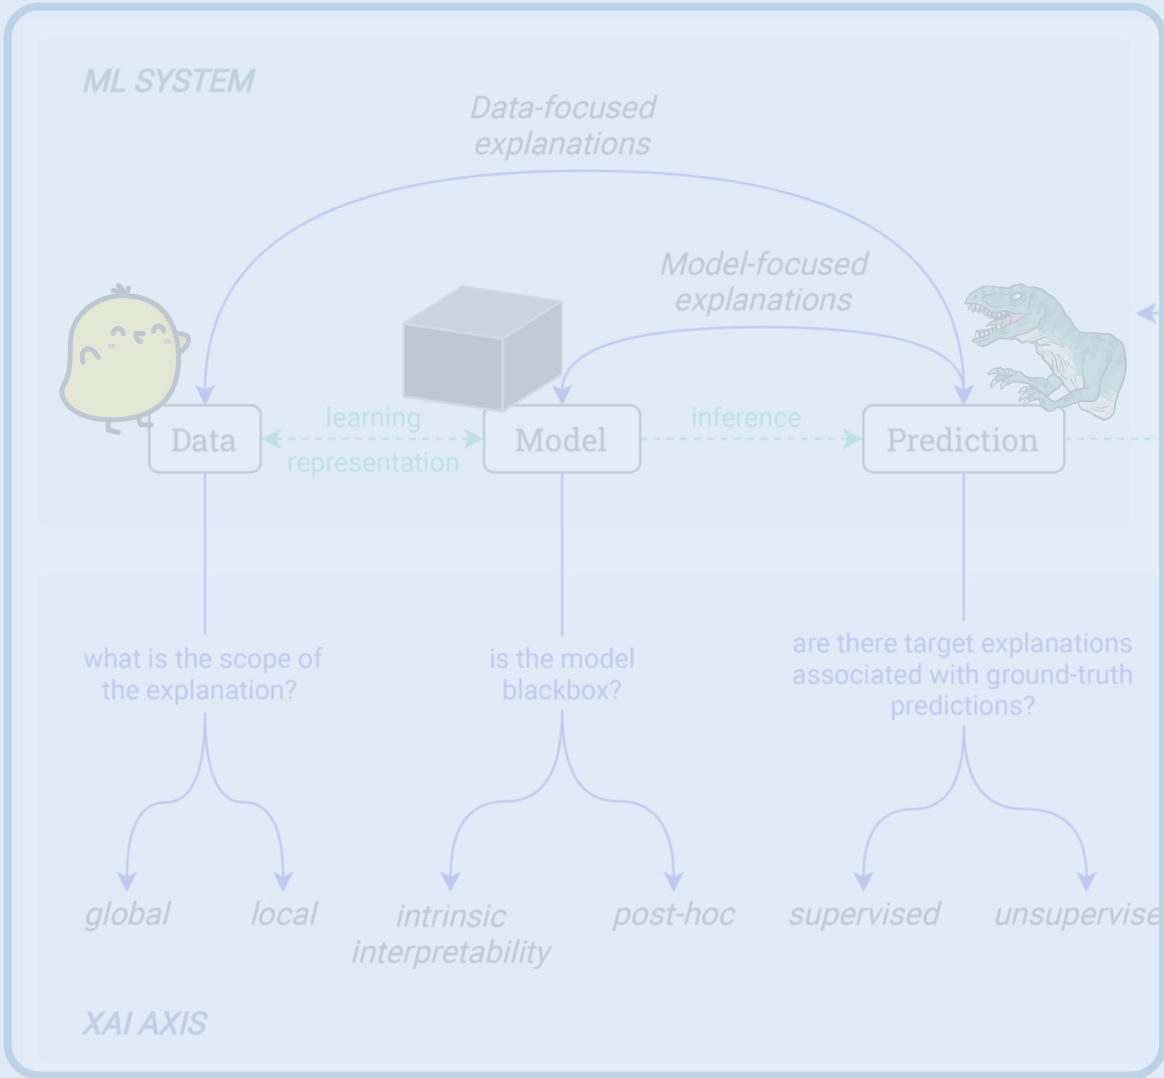


Human-centred XAI approaches

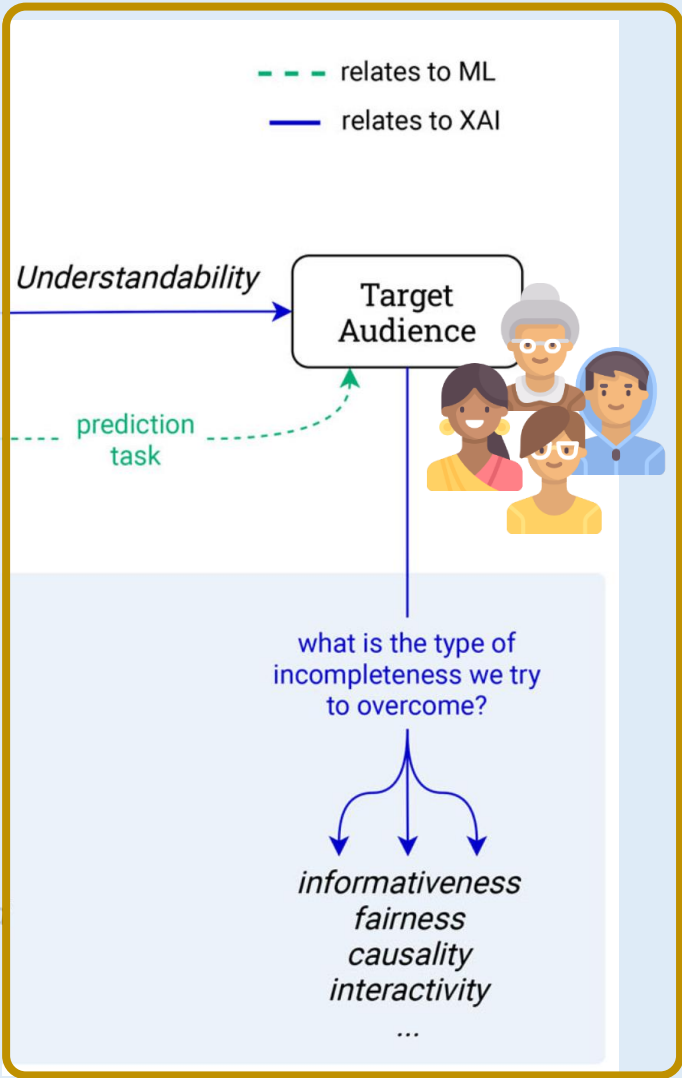


Q. Vera Liao and Kush R. Varshney. 2022. Human-Centered Explainable AI (XAI): From Algorithms to User Experiences. <https://doi.org/10.48550/arXiv.2110.10790>

Algorithmic XAI approaches



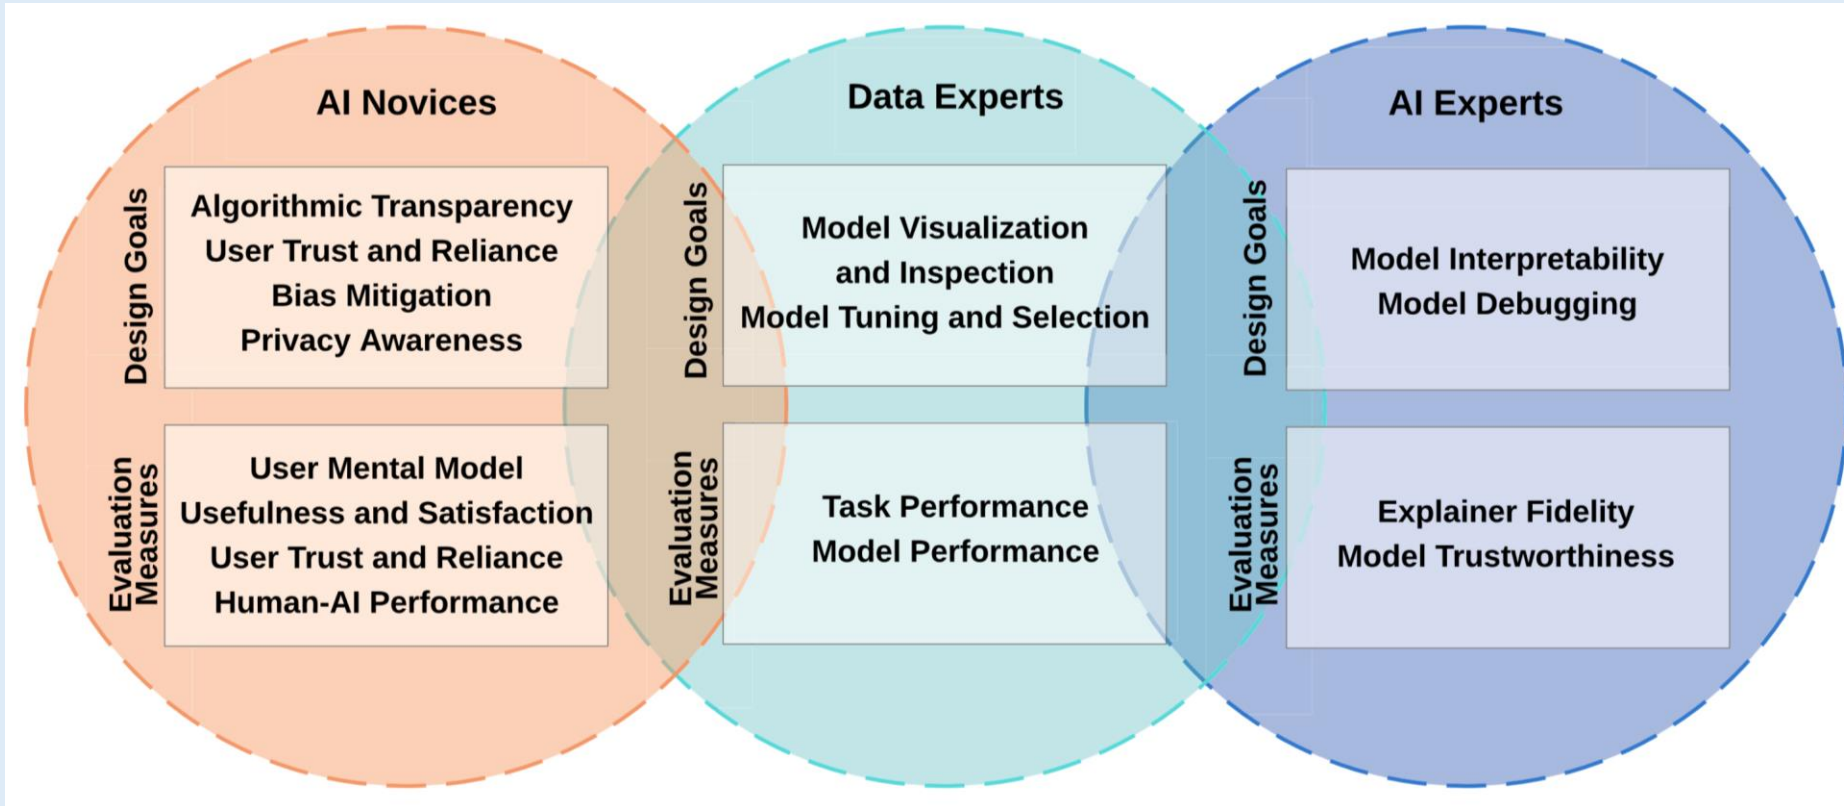
Human-centred XAI approaches



Human-centred XAI approaches

<p>Data</p>	<ul style="list-style-type: none"> • What kind of data was the system trained on? • What is the source of the training data? • How were the labels/ground-truth produced? • What is the sample size of the training data? • What dataset(s) is the system NOT using? • What are the potential limitations/biases of the data? • What is the size, proportion, or distribution of the training data with given feature(s)/feature-value(s)? 	<p>Why</p>	<ul style="list-style-type: none"> • Why/how is this instance given this prediction? • What feature(s) of this instance determine the system's prediction of it? • Why are [instance A and B] given the same prediction?
<p>Output</p>	<ul style="list-style-type: none"> • What kind of output does the system give? • What does the system output mean? • What is the scope of the system's capability? Can it do...? • How is the output used for other system component(s) ? • How should I best utilize the output of the system? • How should the output fit in my workflow? 	<p>Why not</p>	<ul style="list-style-type: none"> • Why is this instance NOT predicted to be [a different outcome Q]? • Why is this instance predicted [P instead of a different outcome Q]? • Why are [instance A and B] given different predictions?
<p>Performance</p>	<ul style="list-style-type: none"> • How accurate/precise/reliable are the predictions? • How often does the system make mistakes? • In what situations is the system likely to be correct/incorrect? • What are the limitations of the system? • What kind of mistakes is the system likely to make? • Is the system's performance good enough for...? 	<p>How to be that (a different prediction)</p>	<ul style="list-style-type: none"> • How should this instance change to get a different prediction Q? • What is the minimum change required for this instance to get a different prediction Q? • How should a given feature change for this instance to get a different prediction Q? • What kind of instance is predicted of [a different outcome Q]?
<p>How (global model-wide explanation)</p>	<ul style="list-style-type: none"> • How does the system make predictions? • What features does the system consider? <ul style="list-style-type: none"> • Is [feature X] used or not used for the predictions? • What is the system's overall logic? <ul style="list-style-type: none"> • How does it weigh different features? • What kind of rules does it follow? • How does [feature X] impact its predictions? • What are the top rules/features that determine its predictions? • What kind of algorithm is used? <ul style="list-style-type: none"> • How were the parameters set? 	<p>How to still be this (the current prediction)</p>	<ul style="list-style-type: none"> • What is the scope of change permitted for this instance to still get the same prediction? • What is the range of value permitted for a given feature for this prediction to stay the same? • What is the necessary feature(s)/feature-value(s) present or absent to guarantee this prediction? • What kind of instance gets the same prediction?
		<p>What If</p>	<ul style="list-style-type: none"> • What would the system predict if this instance changes to...? • What would the system predict if a given feature changes to...? • What would the system predict for [a different instance]?
		<p>Others</p>	<ul style="list-style-type: none"> • How/why will the system change/adapt/improve/drift over time? (change) • Can I, and if so, how do I, improve the system? (improvement) • Why is the system using or not using a given algorithm/feature/rule/dataset? (follow-up) • What does [a machine learning terminology] mean? (terminological) • What are the results of other people using the system? (social)

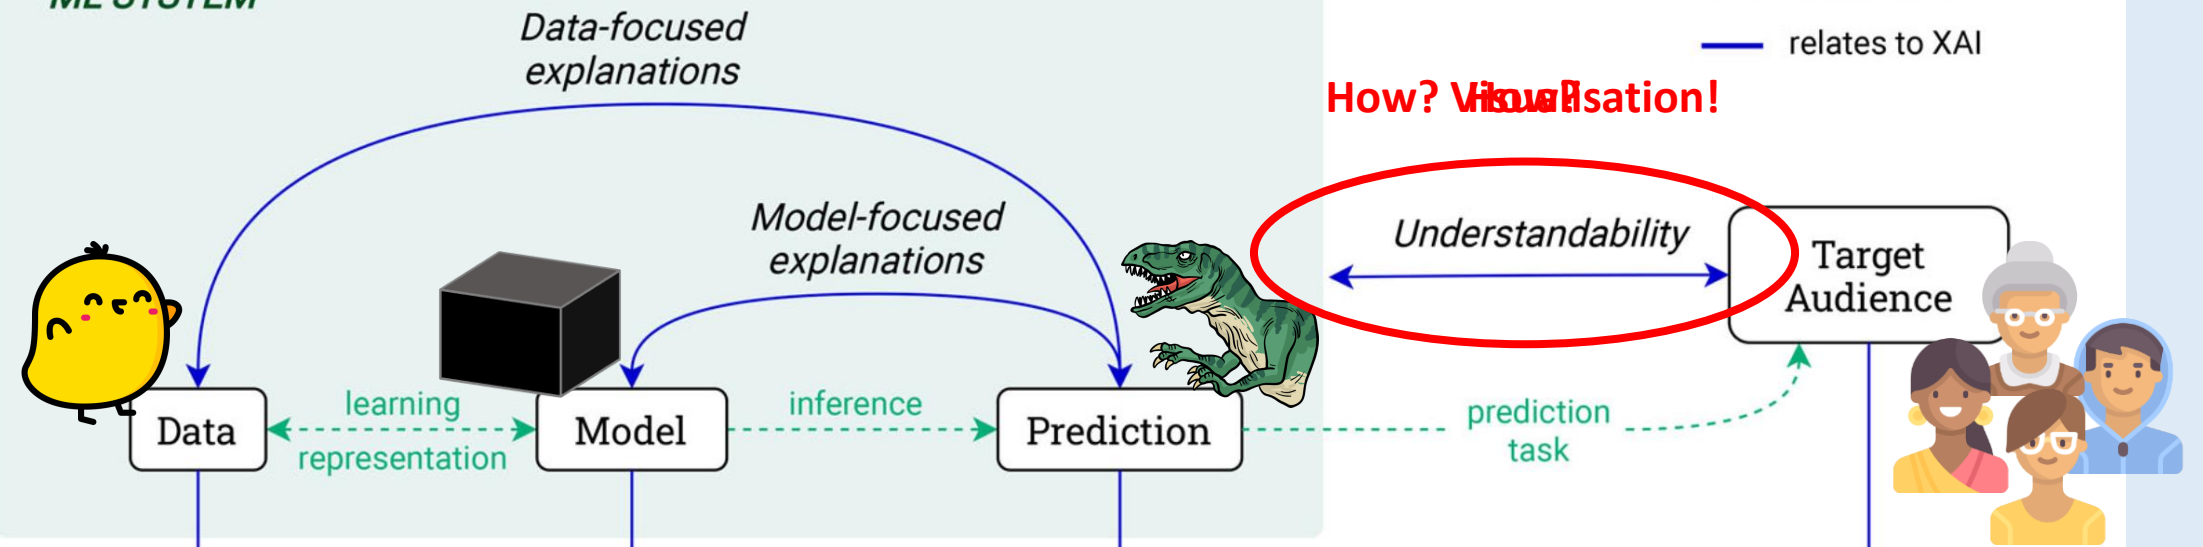
Human-centred XAI approaches



“XAI presents as much of a **design challenge** as an algorithmic challenge” (Q. Vera Liao and Kush R. Varshney, 2022)

ML SYSTEM

--- relates to ML
— relates to XAI



How? Visualisation!

what is the scope of the explanation?

global *local*

is the model blackbox?

intrinsic interpretability *post-hoc*

are there target explanations associated with ground-truth predictions?

supervised *unsupervised*

what is the type of incompleteness we try to overcome?

informativeness
fairness
causality
interactivity
...

XAI AXIS

Visual explanations for AI decisions

Fostering trust with transparency and control



Jeroen Ooge
jeroenooge.com

KU LEUVEN



Katrien Verbert
augment.cs.kuleuven.be

KU LEUVEN

Vraag: Wat is de hoofdstad van de staat Florida?
Antwoord: Tallahassee

Kwaliteit afdelers: Goed

Gevoerdheidsgraad: laag, Gemiddeld, Hoog

Afdelers: Miami, New York, Los Angeles, Limburg, Noorwegen, Zuidpool

visualization
interaction
shepherding
direct explanation

healthcare
visual analytics
algorithms
explainability

opportunities, adapts, can shepherd, often require, how?

gevorderde beginner | Volg mij is dit nu je level voor het onderwerp Hoofdbewerkingen

Welke moeilijkheidsgraad wil je voor de volgende oefeningenreeks?

Heel makkelijk, Makkelijk, Gewoon, Moeilijk, Heel moeilijk

Expert, Bedreven, Competent, Beginner

Start de reeks

Hoe is je nieuw niveau bepaald?

Wiski schat jouw niveau en de moeilijkheid van oefeningen in. Beide veranderen bij het niveau is gestegen na het maken van de reeks oefeningen. Daarna is het nog extra gestegen door je feedback.

Voor reeks, Na feedback

Beginner, Gevorderde beginner, Competent, Bedreven, Expert

Maak meer oefeningen over dit onderwerp | Ga terug naar oefeningsreeks

aangeraden oefening van hetzelfde hoofdstuk

Aangeraden: Oefening 37, Oefening 26, Oefening 21

Waarom deze oefening? Wiski denkt dat jouw huidige niveau past bij dat van deze oefening!

Wiski verwacht dat je 1 of 2 pogingen nodig gaat hebben om oefening 21 juist te maken, gebaseerd op de resultaten van jou en je medeleerlingen.

Aantal pogingen medeleerlingen nodig hadden om oefening 21 juist op te lossen

Maak oefening 21

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht



Patient ID: 3033 | Age: 47 | Risk Factor: High

Region: South | Blood Sugar: 7.8 | Drinking Status: No Alcohol

Intrinsics: 64 | BMI: 33.9 | Walking Status: No Staircase

Gender: Male | Waist Circumference: 112.9 | Physical Activity Level: Low

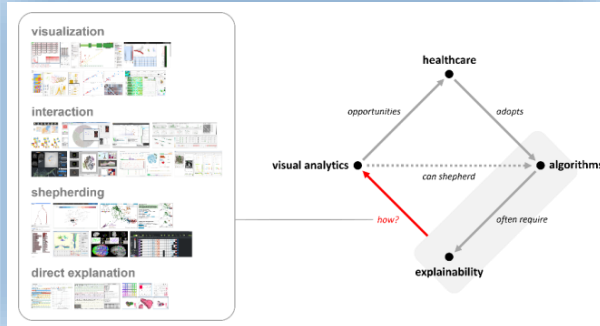
Blood Sugar: 7.5 | Walking Status: 112 | BMI: 33.9

Recommendations to reduce risk: Exercise everyday for 30 min, Before walk increase by 14 cm

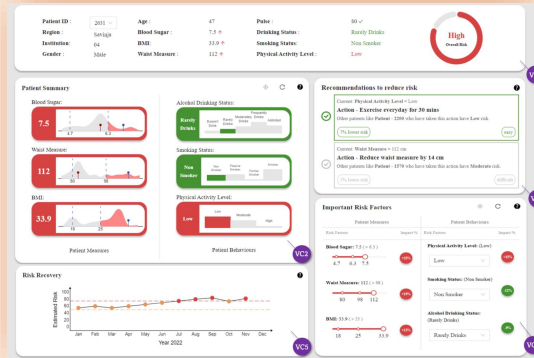
Risk Recovery: 100% | 100%

Explainable AI through visualisation

Visual analytics



Transparency



Maak een aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

- Oefening 37
- Oefening 26
- Oefening 21

Waarom deze oefening? Wiski denkt dat jouw huidige niveau past bij dat van deze oefening! Wiski verwacht dat je 1 of 2 pogingen nodig gaat hebben om oefening 21 juist te maken, gebaseerd op de resultaten van jou en je medeleerlingen.

Aantal pogingen medeleerlingen nodig hadden om oefening 21 juist op te lossen

Maak oefening 21

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht

Control

Hoe is je nieuw niveau bepaald?

Wiski schat jouw niveau en de moeilijkheid van oefeningen in. Beide veranderen bij het oplossen van oefeningen. Je niveau is gestegen na het maken van de reeks oefeningen. Daarna is het nog extra gestegen door je feedback.

Expert

Bedreven

Competent

Gevorderde beginner

Beginner

Voor reeks Na reeks Na feedback

Maak meer oefeningen over dit onderwerp Ga terug naar oefenpagina

How good do you think you are at mathematics?

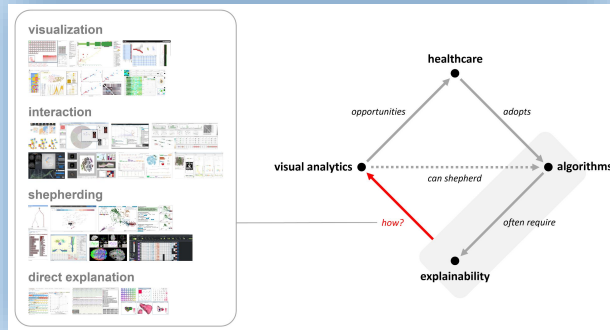
There is no right or wrong answer. Wiski uses your answer to find suitable exercises for you.

- Expert:** mathematics holds no secrets for you.
- Proficient:** you score better than average on mathematics.
- Competent:** you score average on mathematics.
- Advanced beginner:** basic exercises are not a problem for you.
- Novice:** you often have a hard time understanding mathematics.

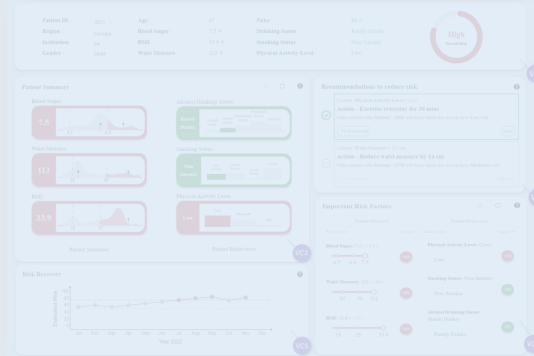
Submit

Explainable AI through visualisation

Visual analytics



Transparency



Maak een aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

- Oefening 27
- Oefening 26
- Oefening 21

Waarom deze oefening? Wiki denkt dat jouw huidige niveau gaat bij dat van deze oefening!

Wiki verwacht dat je 1 of 2 pogingen nodig gaat hebben om oefening 21 juist te maken, gebaseerd op de resultaten van jou en je medeleerlingen.

Aantal pogingen medeleerlingen nodig hadden om oefening 21 juist op te lossen

Maak oefening 21

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht

Control

Hoe is je nieuw niveau bepaald?

Wiki schat jouw niveau en de moeilijkheid van oefeningen in. Beide veranderen bij het oplossen van oefeningen. Je niveau is gestegen na het maken van de reeks oefeningen. Daarna is het nog extra gestegen door je feedback.

Maak meer oefeningen over dit onderwerp

Ga terug naar oefenpagina


How good do you think you are at mathematics?

There is no right or wrong answer. Wiki uses your answer to find suitable exercises for you.


- Expert:** mathematics holds no secrets for you.
- Proficient:** you score better than average on mathematics.
- Competent:** you score average on mathematics.
- Advanced beginner:** basic exercises are not a problem for you.
- Novice:** you often have a hard time understanding mathematics.

Submit


visualization



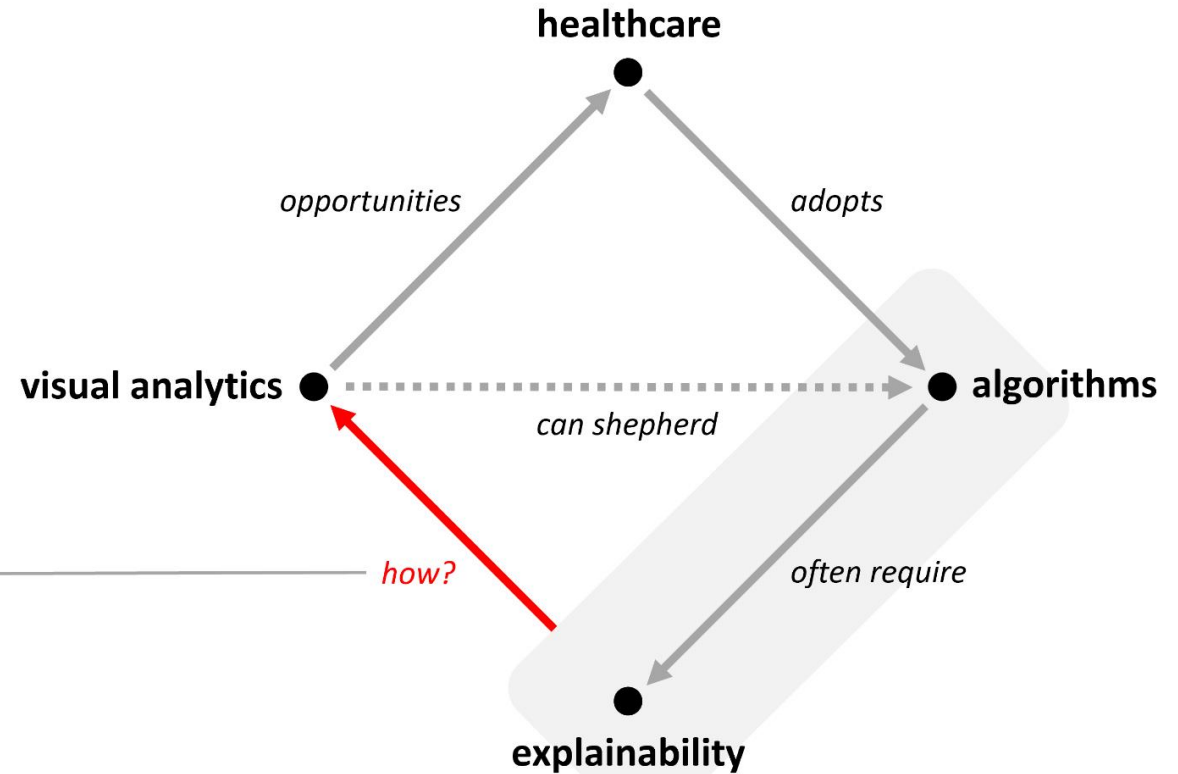
interaction



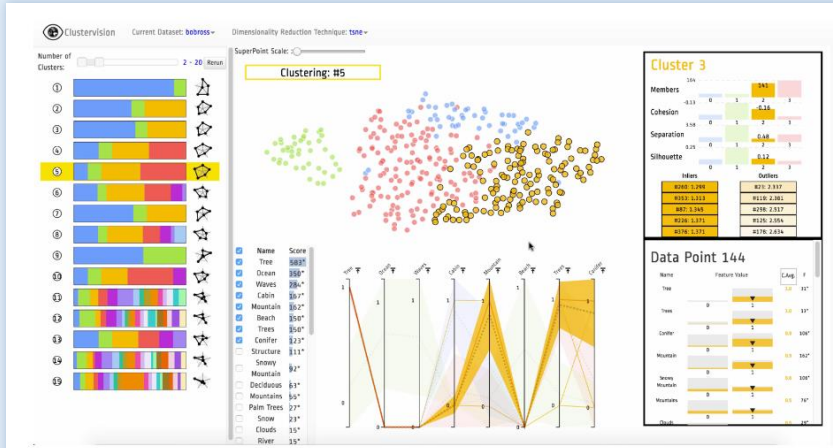
shepherding



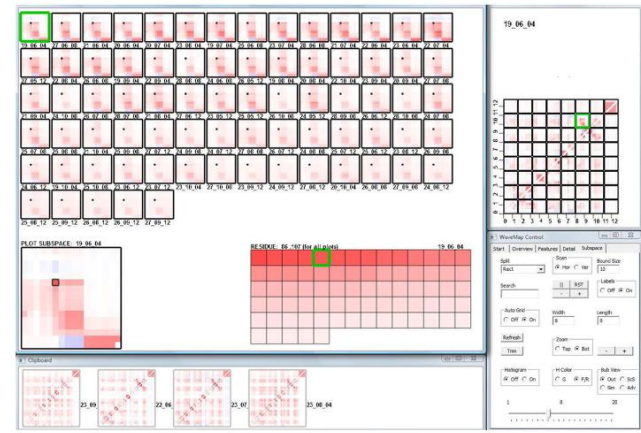
direct explanation



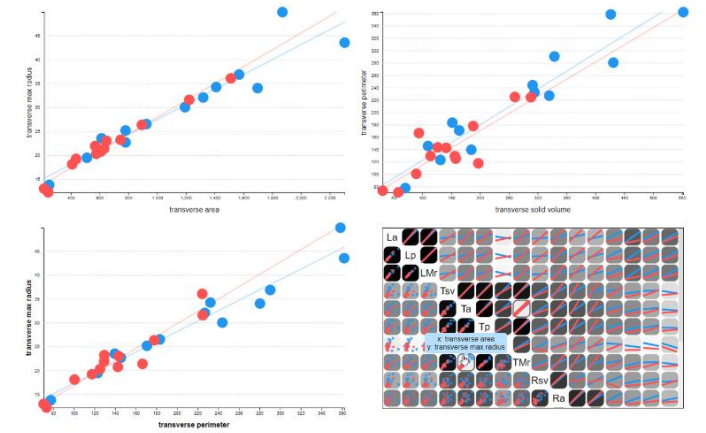
Visualisation



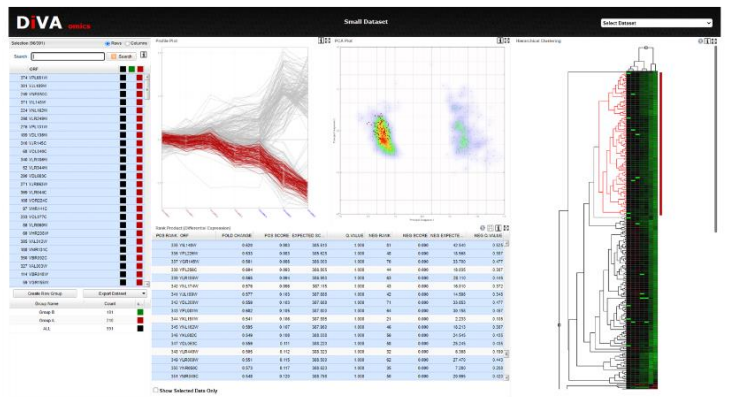
(a) Clustering




(b) Similarity



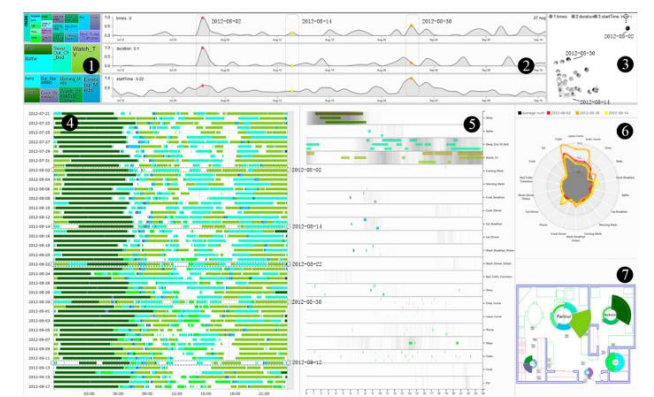
(c) Classical statistics



(d) Dimension reduction

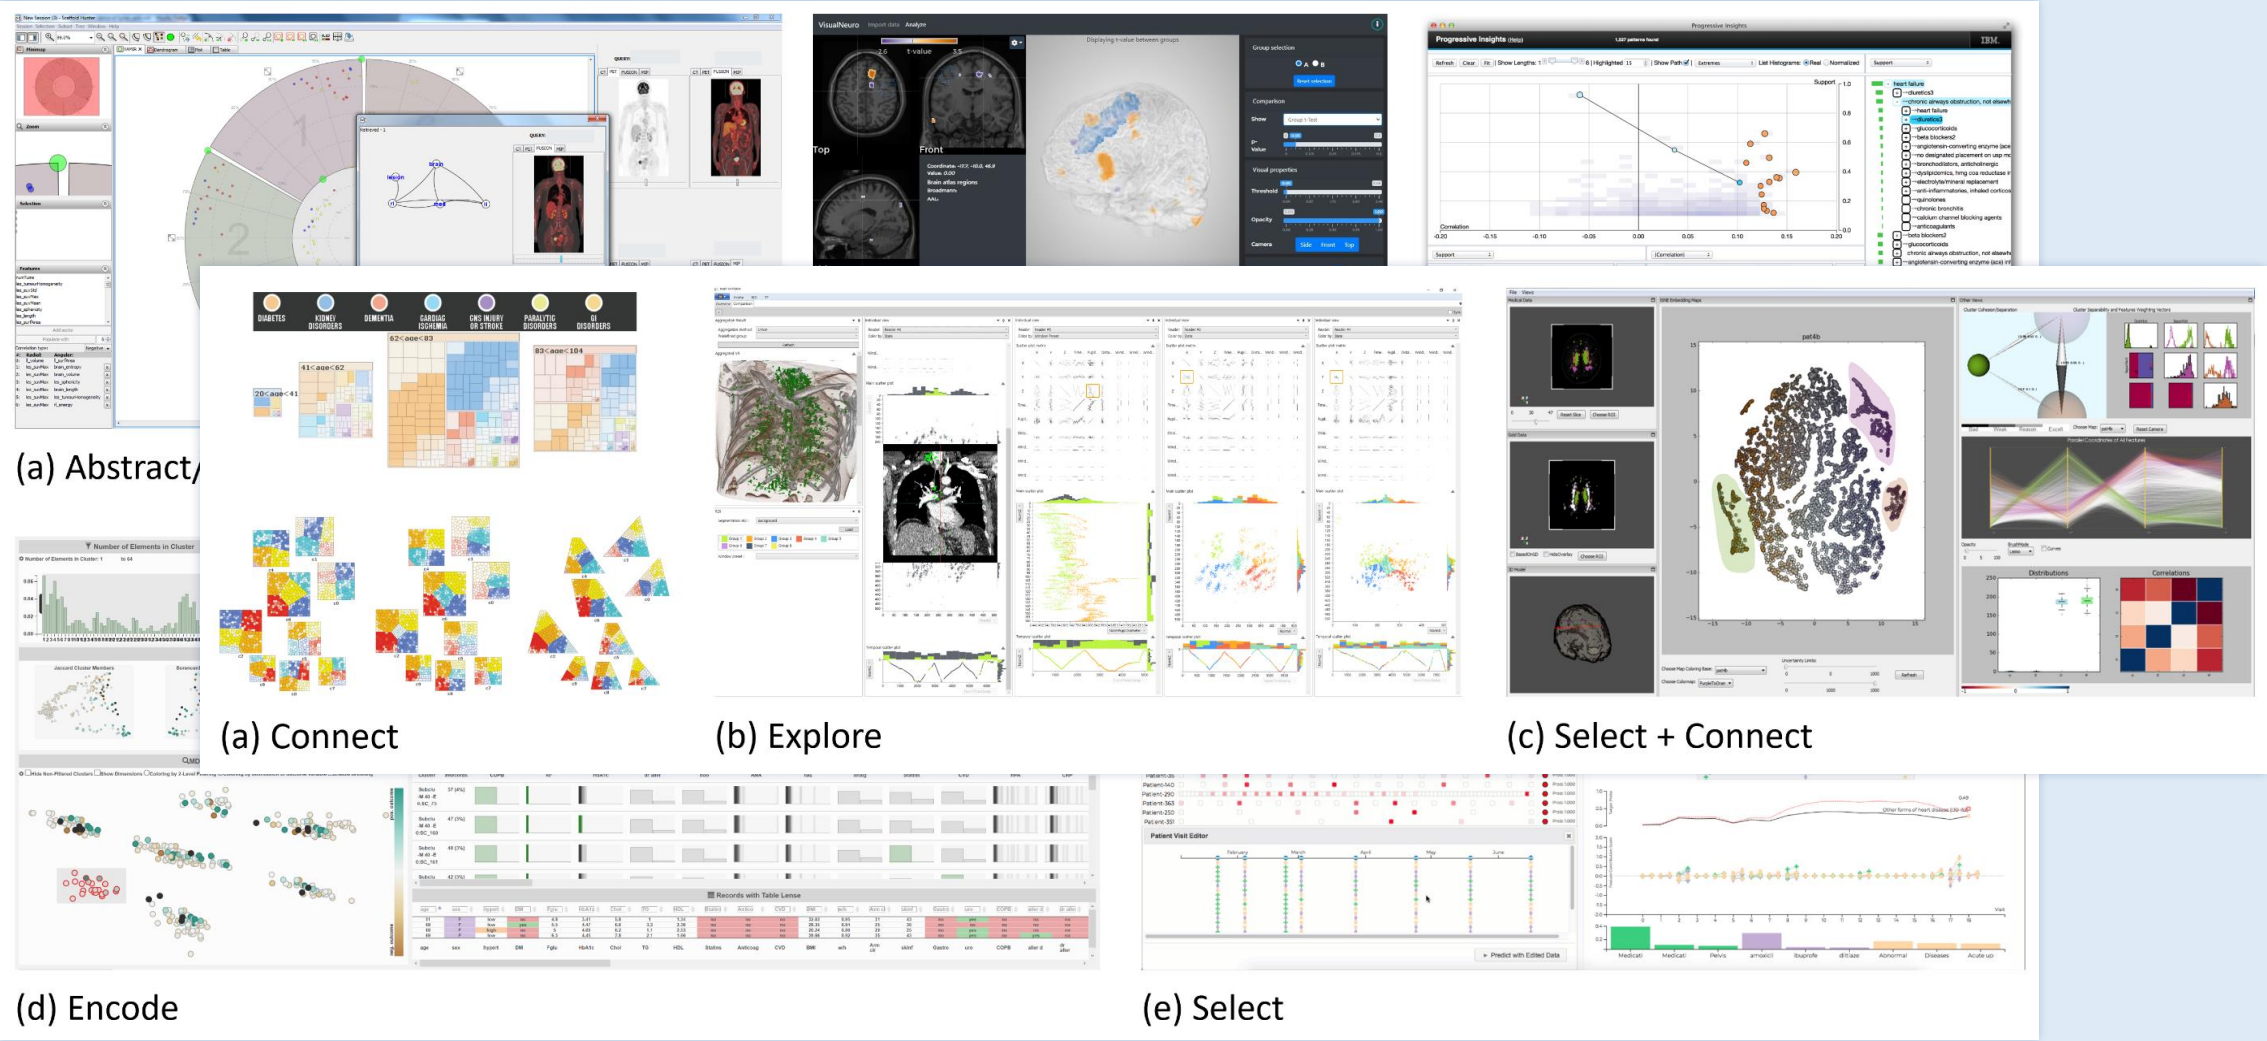


(e) Data mining




(f) Anomaly detection

Interaction




Jeroen Ooge, Gregor Stiglic, and Katrien Verbert. 2022. Explaining artificial intelligence with visual analytics in healthcare. *WIREs Data Mining and Knowledge Discovery* 12, 1: e1427. <https://doi.org/10.1002/widm.1427>

Shepherding



(a) Configuration window separate from visualization




(c) Fixed settings panel, automatically rerun algorithm




(e) Automatically rerun algorithm when input features change

Level 2
semi-interactive


Level 3
tight integration



(b) Fixed settings panel, manually rerun algorithm

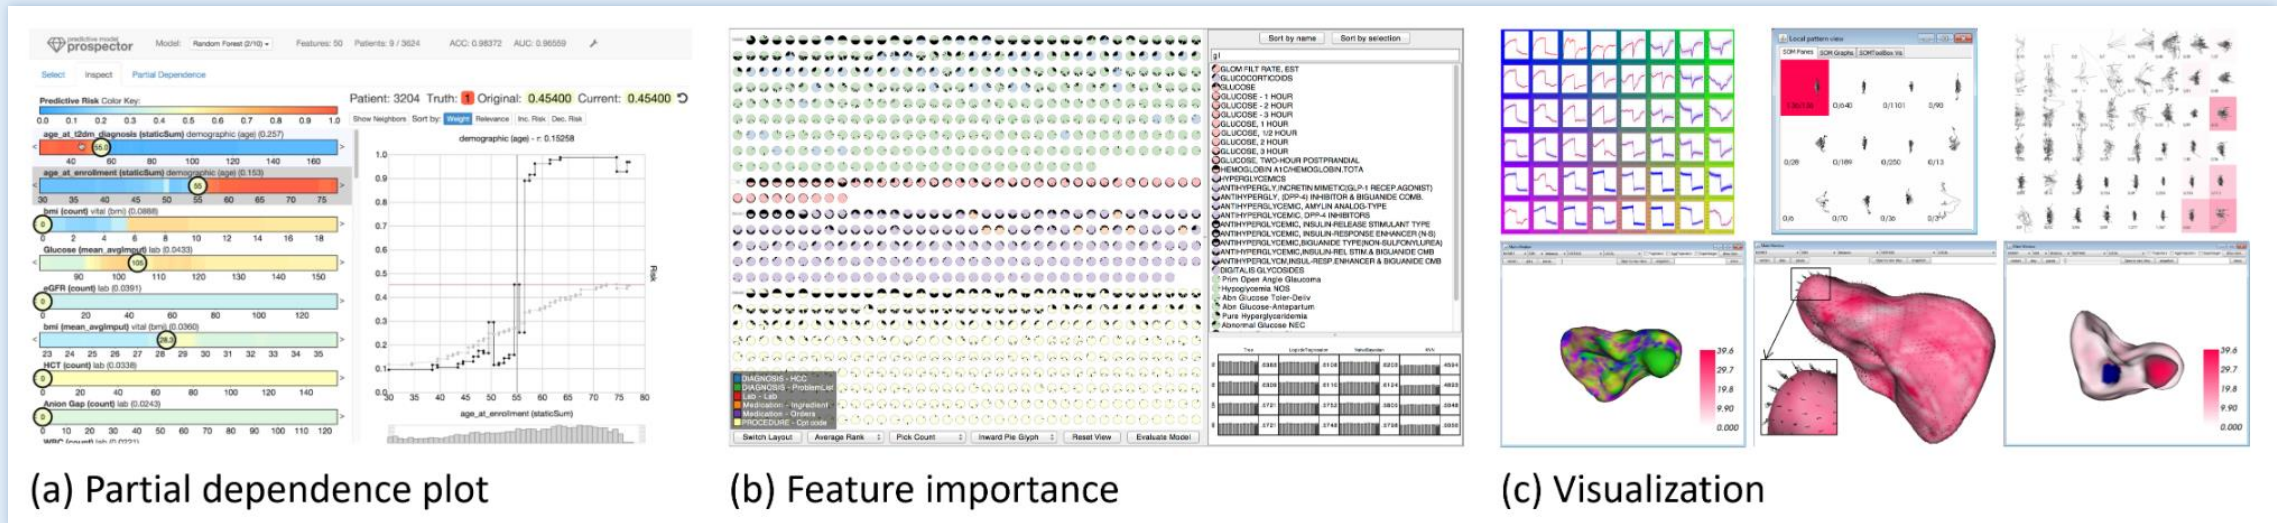


(d) Configuration in visualization interface, automatically rerun algorithm



(f) Interact with visualization



Direct explanation



Visual analytics systems targeting **laypeople**, supporting shepherding, or containing direct explanations are rare!



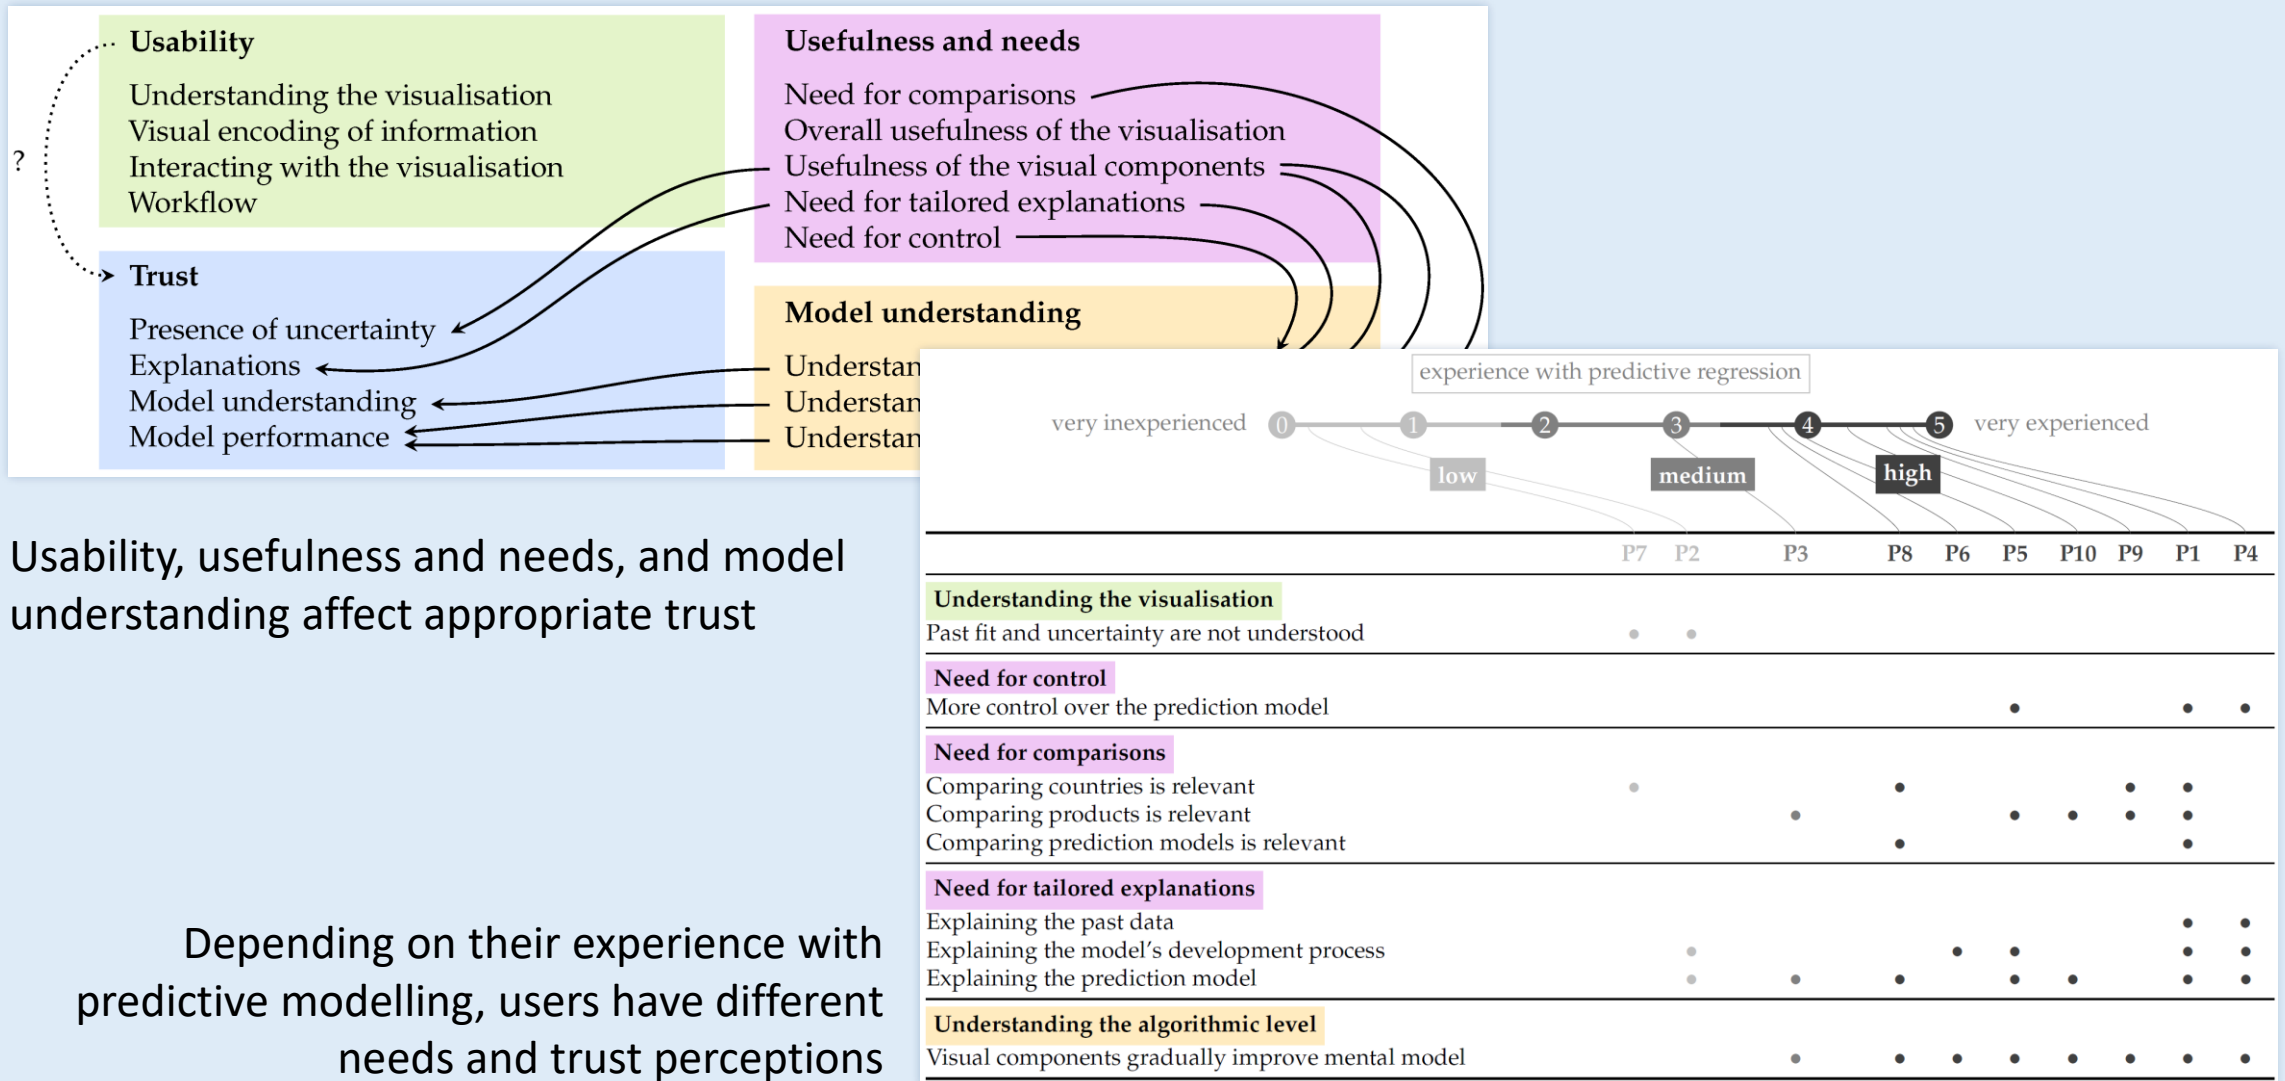
Jeroen Ooge, Gregor Stiglic, and Katrien Verbert. 2022. Explaining artificial intelligence with visual analytics in healthcare. *WIREs Data Mining and Knowledge Discovery* 12, 1: e1427. <https://doi.org/10.1002/widm.1427>



Experts have **different trust evolutions** when using an unknown prediction model



Jeroen Ooge and Katrien Verbert. 2021. Trust in Prediction Models: a Mixed-Methods Pilot Study on the Impact of Domain Expertise. In 2021 IEEE Workshop on TRust and EXPertise in Visual Analytics (TRES), 8–13. <https://doi.org/10.1109/TRES53765.2021.00007>

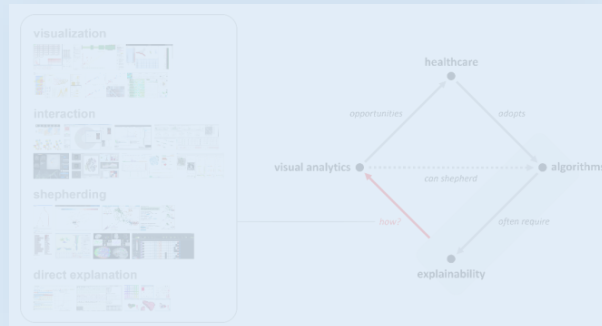


Usability, usefulness and needs, and model understanding affect appropriate trust

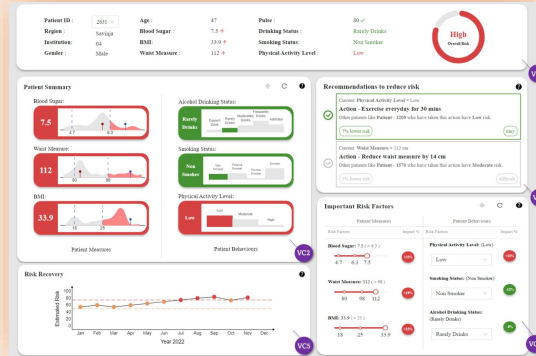
Depending on their experience with predictive modelling, users have different needs and trust perceptions

Explainable AI through visualisation

Visual analytics



Transparency




Maak een aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

- Oefening 37
- Oefening 26
- Oefening 21

Waarom deze oefening? Wiski denkt dat jouw huidige niveau past bij dat van deze oefening!
Wiski verwacht dat je **1 of 2 pogingen** nodig gaat hebben om oefening 21 juist te maken, gebaseerd op de resultaten van jou en je medeleerlingen.



[Maak oefening 21](#)

... of kies zelf je volgende oefening

[Naar het oefeningenoverzicht](#)

Control

Hoe is je nieuw niveau bepaald?

Wiski schat jouw niveau en de moeilijkheid van oefeningen in. Beide veranderen bij het oplossen van oefeningen. Je niveau is gestegen na het maken van de reeks oefeningen. Daarna is het nog extra gestegen door je feedback.



[Maak meer oefeningen over dit onderwerp](#)

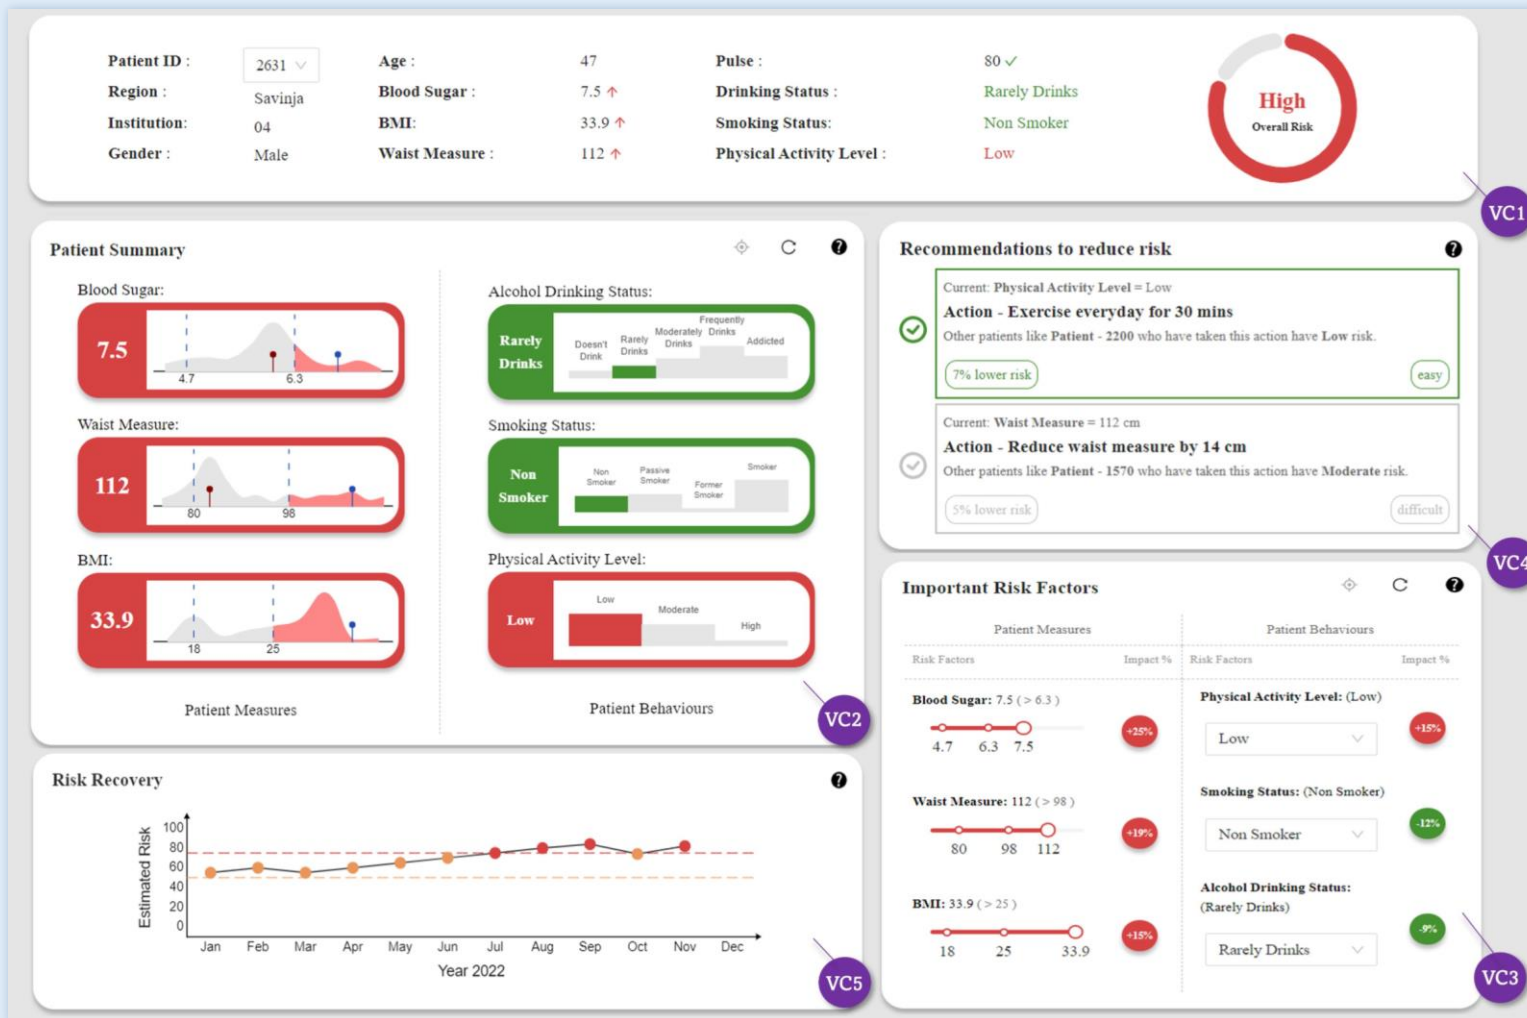
[Ga terug naar oefenpagina](#)

How good do you think you are at mathematics?

There is no right or wrong answer. Wiski uses your answer to find suitable exercises for you.

- Expert: mathematics holds no secrets for you.
- Proficient: you score better than average on mathematics.
- Competent: you score average on mathematics.
- Advanced beginner: basic exercises are not a problem for you.
- Novice: you often have a hard time understanding mathematics.

[Submit](#)



Participants preferred data-centric explanations that provide local explanations with a global overview over other methods



Aditya Bhattacharya, Jeroen Ooge, Gregor Stiglic, and Katrien Verbert. 2023. Directive Explanations for Monitoring the Risk of Diabetes Onset: Introducing Directive Data-Centric Explanations and Combinations to Support What-If Explorations. In *Proceedings of the 28th International Conference on Intelligent User Interfaces (IUI '23)*, 204–219. <https://doi.org/10.1145/3581641.3584075>



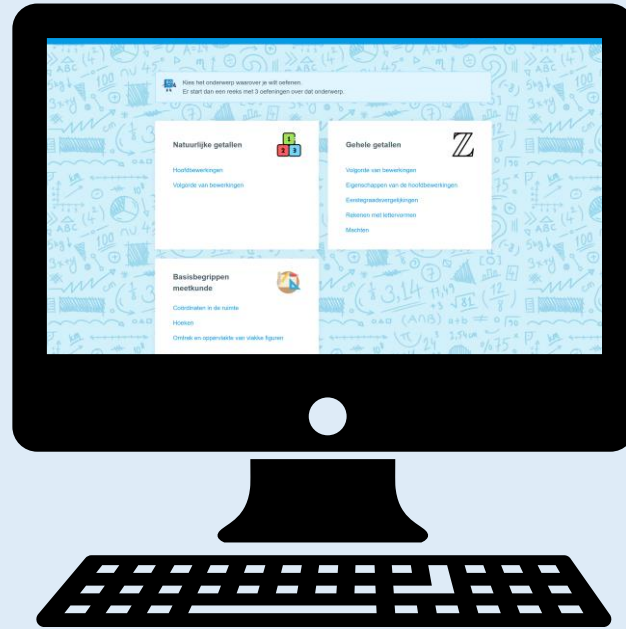

 expert




 competent




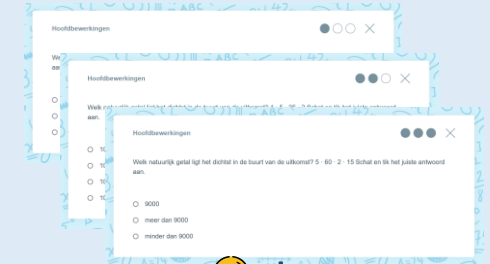
 beginner



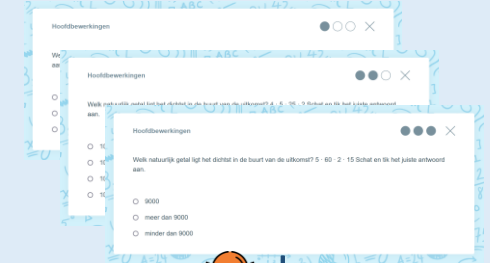
Transparency: why these exercises?




 because ...



 because ...



 because ...

Maak een aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

 Oefening 37


 Oefening 26

 Oefening 21

Waarom deze oefening? Wiski denkt dat jouw huidig niveau past bij dat van deze oefening!

Wiski verwacht dat je **1 of 2 pogingen** nodig gaat hebben om oefening 21 juist te maken, gebaseerd op de resultaten van jou en je medeleerlingen.

Aantal pogingen medeleerlingen nodig hadden om oefening 21 juist op te lossen



Maak oefening 21

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht

Textual explanation

Visual explanation

Maak een aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

 Oefening 27

 Oefening 40

 Oefening 45

Waarom deze oefening?

Oefening 27 is aangeraden omdat het algoritme van Wiski dat zo heeft berekend.



Maak oefening 27

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht

Placebo explanation

Maak een aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

 Oefening 27

 Oefening 40

 Oefening 45

Wiski raadt de volgende oefening aan

GOOD JOB 


Maak oefening 27

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht

No explanation

Randomised controlled experiment



Final questionnaire: trust

Competence

- Q1 Wiski is like an expert (for example, a teacher) for recommending exercises.
- Q2 Wiski has the expertise (knowledge) to estimate my math level.
- Q3 Wiski can estimate my math level.
- Q4 Wiski understands the difficulty level of math exercises well.
- Q5 Wiski takes my math level into account when recommending exercises.

Benevolence

- Q6 Wiski prioritizes that I improve in math.
- Q7 Wiski recommends exercises so that I improve in math.
- Q8 Wiski wants to estimate my math level well.

Integrity

- Q9 Wiski recommends exercises as correctly as possible.
- Q10 Wiski is honest.
- Q11 Wiski makes integrous recommendations.

Trust (one-dimensional)

- Q12 I trust Wiski to recommend me math exercises.

Intention to return

- Q13 If I want to solve math exercises again, I will choose Wiski.
- Q14 If I want to be recommended math exercises again, I will choose Wiski.

Perceived transparency

- Q15 I find that Wiski gives enough explanation as to why an exercise is recommended.

Trusting beliefs


Multidimensional trust

Trust effects

PLACEBO



1D trust




MD trust



REAL



1D trust



MD trust



Maak een aangeraden oefening van hetzelfde hoofdstuk


Aangeraden

-  Oefening 37
-  Oefening 26
-  Oefening 21

Waarom deze oefening? Wiski denkt dat jouw huidig niveau past bij dat van deze oefening!

Wiski verwacht dat je **1 of 2 pogingen** nodig gaat hebben om oefening 21 juist te maken, gebaseerd op de resultaten van jou en je medeleerlingen.

Aantal pogingen medeleerlingen nodig hadden om oefening 21 juist op te lossen



Maak oefening 21

Wiski raadt je de volgende oefening

aan te maken op basis van het oefeningenoverzicht

Maak een aangeraden oefening van hetzelfde hoofdstuk

- Aangeraden
-  Oefening 27
 -  Oefening 40
 -  Oefening 45

Wiski raadt de volgende oefening aan



Maak oefening 27

... of kies zelf je volgende oefening

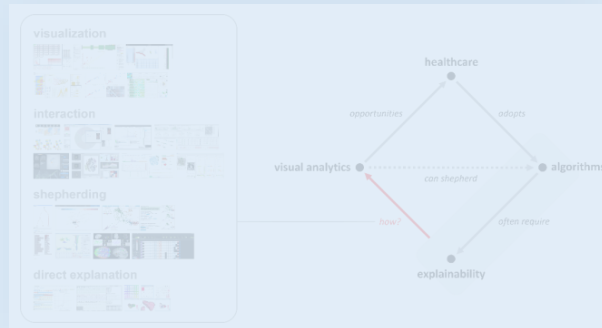
Naar het oefeningenoverzicht

Visual explanations can increase trust may not be the most important factor for building it

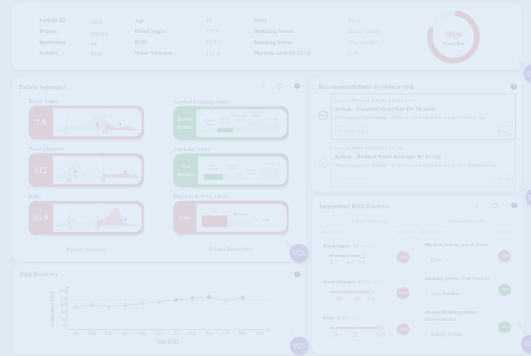


Explainable AI through visualisation

Visual analytics



Transparency



Maak een aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

- Oefening 27
- Oefening 26
- Oefening 21

Waarom deze oefening? Wiski denkt dat jouw huidige niveau gaat bij dat van deze oefening! Wiski verwacht dat je 1 of 2 pagijnen nodig gaat hebben om oefening 21 juist te maken, gebaseerd op de resultaten van jou en je medeleerlingen.

Aantal pogingen medeleerlingen nodig hadden om oefening 21 juist te maken

The bar chart shows the number of attempts needed by other students to solve exercise 21 correctly. The x-axis represents the number of attempts (1, 2, 3, 4), and the y-axis represents the number of students. The highest number of students needed 1 attempt, followed by 2 attempts, and then 3 and 4 attempts.

Maak oefening 21

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht

Control

Hoe is je nieuw niveau bepaald?

Wiski schat jouw niveau en de moeilijkheid van oefeningen in. Beide veranderen bij het oplossen van oefeningen. Je niveau is gestegen na het maken van de reeks oefeningen. Daarna is het nog extra gestegen door je feedback.

The graph shows the relationship between level and feedback. The x-axis represents the level (Beginner, Gevorderde beginner, Competent, Bedreven, Expert) and the y-axis represents the level. A green line shows the level increasing from 'Voor reeks' to 'Na reeks' to 'Na feedback'. The level is highest for 'Na feedback'.

Maak meer oefeningen over dit onderwerp

Ga terug naar oefenpagina

How good do you think you are at mathematics?

There is no right or wrong answer. Wiski uses your answer to find suitable exercises for you.

- Expert:** mathematics holds no secrets for you.
- Proficient:** you score better than average on mathematics.
- Competent:** you score average on mathematics.
- Advanced beginner:** basic exercises are not a problem for you.
- Novice:** you often have a hard time understanding mathematics.

Submit



 expert

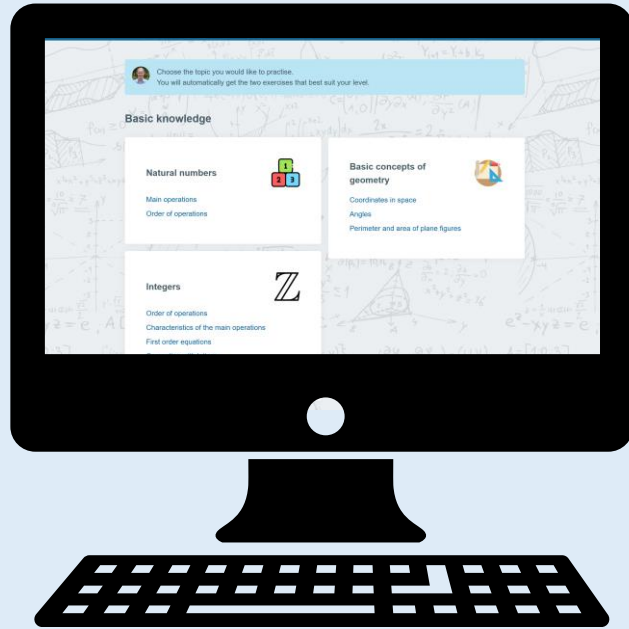


 competent



 beginner

$$\begin{matrix} \bullet & \sqrt{x} & & + \\ & \int dx & & \\ \pi & & & \blacktriangle \end{matrix}$$



Control: I want other exercises





WISKI

Oefen wiskundeleerstof van de middelbare school

Inloggen

Registreren



Leren op maat

Honderden meerkeuzevragen over 42 onderwerpen uit de leerstof van het middelbaar: kies zelf wat je wilt oefenen.



Oefenen op jouw niveau

Niet elke leerling heeft hetzelfde niveau. Wiski helpt jou om de beste oefeningen te maken voor jouw niveau.




Geef feedback

Ben je het niet helemaal eens met Wiski? Wiski zal naar je luisteren en er rekening mee houden.

Randomised controlled experiment



76
adolescents



Wiski would like additional information from you

You solved a complete series of recommended exercises, congratulations! For the next series, you can give Wiski additional information so that Wiski knows better how you feel.

What difficulty of exercises would you like?

Easier Similar Harder

Submit feedback



How is your new level determined?

Wiski estimates your level and the difficulty of exercises. Both change when solving exercises. Your level remained similar after solving the exercise series. Then, it increased even further because of your feedback.

Expert
Proficient
Competent
Advanced beginner
Novice

Before series After series After feedback

Solve more exercises on this topic Return to exercise page



NONE

Wiski would like additional information from you

You solved a complete series of recommended exercises, congratulations! For the next series, you can give Wiski additional information so that Wiski knows better how you feel.

What difficulty of exercises would you like?

Easier Similar Harder

Submit feedback



How is your new level determined?

Wiski estimates your level and the difficulty of exercises. Both change when solving exercises. Your level remained similar after solving the exercise series. Then, it increased even further because of your feedback.

Expert
Proficient
Competent
Advanced beginner
Novice

Before series After series After feedback

Solve more exercises on this topic Return to exercise page



CONTROL

Wiski would like additional information from you

You solved a complete series of recommended exercises, congratulations! For the next series, you can give Wiski additional information so that Wiski knows better how you feel.

What difficulty of exercises would you like?

Easier Similar Harder

Submit feedback



How is your new level determined?

Wiski estimates your level and the difficulty of exercises. Both change when solving exercises. Your level remained similar after solving the exercise series. Then, it increased even further because of your feedback.

Expert
Proficient
Competent
Advanced beginner
Novice

Before series After series After feedback

Solve more exercises on this topic Return to exercise page



CONTROL
+ IMPACT

Final questionnaire: control

Control

- Q18 I feel in control of telling Wiski what I want.
- Q19 I don't feel in control of telling Wiski what I want.
- Q20 I don't feel in control of specifying and changing my preferences.
- Q21 Wiski seems to control my decision process rather than I do.

Preference elicitation

- Q22 Wiski provides an adequate way for me to express my preferences.
- Q23 I found it easy to tell Wiski about my preferences.
- Q24 It is easy to learn to tell Wiski what I like.
- Q25 It required too much effort to tell Wiski what I like.

Preference revision

- Q26 Wiski provides an adequate way for me to revise my preferences.
- Q27 I found it easy to make Wiski recommend different things.
- Q28 It is easy to train Wiski to update my preferences.
- Q29 I found it easy to alter the recommended exercises.
- Q30 It is easy for me to inform Wiski if I dislike/like recommended exercises.
- Q31 It is easy for me to get a new set of recommended exercises.

7-point Likert-type questions + open comments

Final questionnaire: trust

Competence

- Q1 Wiski is like an expert (for example, a teacher) for recommending exercises.
- Q2 Wiski has the expertise (knowledge) to estimate my maths level.
- Q3 Wiski can estimate my maths level.
- Q4 Wiski understands the difficulty level of maths exercises.
- Q5 Wiski takes my maths level into account when recommending exercises.

Benevolence

- Q6 Wiski prioritises that I improve in maths.
- Q7 Wiski recommends exercises so that I improve in maths.
- Q8 Wiski wants to estimate my maths level well.

Integrity

- Q9 Wiski recommends exercises as correctly as possible.
- Q10 Wiski is honest.
- Q11 Wiski makes integrous recommendations.

Trust (one-dimensional)

- Q12 I trust Wiski to recommend me maths exercises.

Intention to return

- Q13 If I want to solve maths exercises again, I will choose Wiski.
- Q14 If I want to be recommended maths exercises again, I will choose Wiski.

Transparency

- Q15 I understood why the exercises were recommended.
- Q16 Wiski helps me understand why the exercises were recommended.
- Q17 Wiski explains why the exercises are recommended.

Trusting beliefs

Multidimensional trust

What are the effects of the **control mechanism**?

	NONE vs. CONTROL	NONE vs. CONTROL+IMPACT	CONTROL vs. CONTROL+IMPACT
Benevolence	0.16 ($p = 0.263$)	0.61 ($p = 0.011$)	0.45 ($p = 0.035$)
Trusting beliefs	-0.01 ($p = 0.529$)	0.38 ($p = 0.042$)	0.40 ($p = 0.030$)
Transparency	0.29 ($p = 0.068$)	1.04 ($p = 0.000$ **)	0.74 ($p = 0.002$ *)
One-dimens. trust	0.00 ($p = 0.504$)	0.78 ($p = 0.017$)	0.78 ($p = 0.020$)
Multidimens. trust	0.15 ($p = 0.207$)	0.55 ($p = 0.009$ *)	0.40 ($p = 0.039$)
Preference revision	0.33 ($p = 0.080$)	0.43 ($p = 0.030$)	0.10 ($p = 0.325$)

* $p < 0.01$, ** $p < 0.001$, non-significant results ($p \geq 0.5$) are greyed out

Wiski would like additional information from you

You solved a complete series of recommended exercises, congratulations! For the next series, you can give Wiski additional information so that Wiski knows better how you feel.

What difficulty of exercises would you like?

Easier Similar Harder

Trust in the e-learning platform



Reflection on own mastery



Reflection on recommendations



What are the effects of visualising the control's impact?

	NONE vs. CONTROL	NONE vs. CONTROL+IMPACT	CONTROL vs. CONTROL+IMPACT
Benevolence	0.16 ($p = 0.263$)	0.61 ($p = 0.011$)	0.45 ($p = 0.035$)
Trusting beliefs	-0.01 ($p = 0.529$)	0.38 ($p = 0.042$)	0.40 ($p = 0.030$)
Transparency	0.29 ($p = 0.068$)	1.04 ($p = 0.000$ **)	0.74 ($p = 0.002$ *)
One-dimens. trust	0.00 ($p = 0.504$)	0.78 ($p = 0.017$)	0.78 ($p = 0.020$)
Multidimens. trust	0.15 ($p = 0.207$)	0.55 ($p = 0.009$ *)	0.40 ($p = 0.039$)
Preference revision	0.33 ($p = 0.080$)	0.43 ($p = 0.030$)	0.10 ($p = 0.325$)

* $p < 0.01$, ** $p < 0.001$, non-significant results ($p \geq 0.5$) are greyed out

The screenshot shows two overlapping windows from the Wiski platform. The top window, titled "Wiski would like additional information from you", congratulates the user and asks for feedback on exercise difficulty, with a slider ranging from "Similar" to "Harder". The bottom window, titled "How is your new level determined?", explains that the user's level is based on exercise difficulty and provides a line graph showing the user's level progression from "Before series" to "After series". The graph shows a steady increase from the "Advanced beginner" level to the "Competent" level. A "Submit feedback" button is visible in the top window, and "Solve more exercises on this topic" and "Return to exercise page" buttons are in the bottom window.

Trust in the e-learning platform



Reflection



Understanding through model inspection?



How is your new level determined?

Wiski estimates your level and the difficulty of exercises. Both change when solving exercises.

Your level remained similar after solving the exercise series.
Then, it increased even further because of your feedback.

Expert

Proficient

Competent

Advanced beginner

Novice

Before series After series After feedback

[Solve more exercises on this topic](#)

How good do you think you are at mathematics?

There is no right or wrong answer. Wiski uses your answer to find suitable exercises for you.

- Expert:** mathematics holds no secrets for you.
- Proficient:** you score better than average on mathematics.
- Competent:** you score average on mathematics.
- Advanced beginner:** basic exercises are not a problem for you.
- Novice:** you often have a hard time understanding mathematics.

[Submit](#)

Wiski would like additional information from you

You solved a complete series of recommended exercises, congratulations! For the next series, you can give Wiski additional information so that Wiski knows better how you feel.

What difficulty of exercises would you like?

Easier Similar Harder

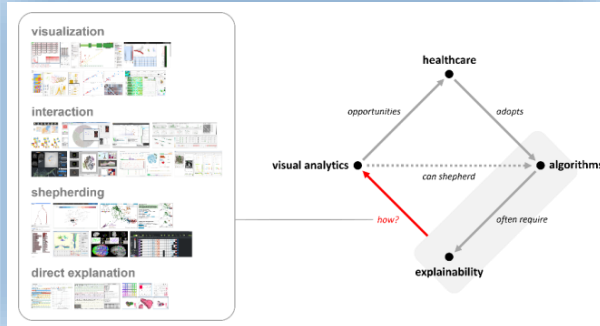
[Submit feedback](#)

Control mechanisms do not necessarily increase trust; showing the impact of control is essential

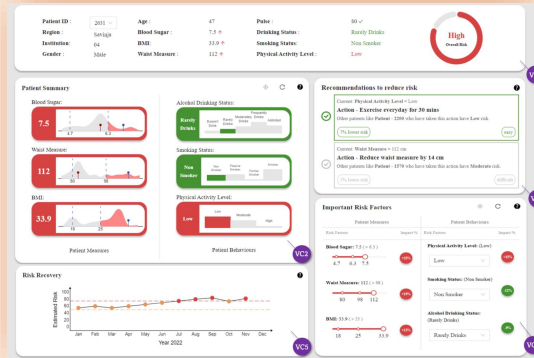


Explainable AI through visualisation

Visual analytics



Transparency



Maak een aangeraden oefening van hetzelfde hoofdstuk

Aangeraden

- Oefening 37
- Oefening 26
- Oefening 21

Waarom deze oefening? Wiski denkt dat jouw huidige niveau past bij dat van deze oefening! Wiski verwacht dat je 1 of 2 pogingen nodig gaat hebben om oefening 21 juist te maken, gebaseerd op de resultaten van jou en je medeleerlingen.

Aantal pogingen medeleerlingen nodig hadden om oefening 21 juist op te lossen

Maak oefening 21

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht

Control

Hoe is je nieuw niveau bepaald?

Wiski schat jouw niveau en de moeilijkheid van oefeningen in. Beide veranderen bij het oplossen van oefeningen. Je niveau is gestegen na het maken van de reeks oefeningen. Daarna is het nog extra gestegen door je feedback.

Expert

Bedreven

Competent

Gevorderde beginner

Beginner

Voor reeks Na reeks Na feedback

Maak meer oefeningen over dit onderwerp Ga terug naar oefenpagina

How good do you think you are at mathematics?

There is no right or wrong answer. Wiski uses your answer to find suitable exercises for you.

- Expert:** mathematics holds no secrets for you.
- Proficient:** you score better than average on mathematics.
- Competent:** you score average on mathematics.
- Advanced beginner:** basic exercises are not a problem for you.
- Novice:** you often have a hard time understanding mathematics.

Submit

Visual explanations for AI decisions

Fostering trust with transparency and control



Jeroen Ooge
jeroenooge.com

KU LEUVEN



Katrien Verbert
augment.cs.kuleuven.be

KU LEUVEN

Vraag: Wat is de hoofdstad van de staat Florida?

Antwoord: Tallahassee

Kwaliteit antwoorden: Goed


Keuzelijst: Miami, New York, Los Angeles, Limburg, Noorwegen, Zuidpool

visualization

interaction

shepherding

direct explanation



gevorderde beginner

Volgens mij is dit nu je level voor het onderwerp Hoofdbewerkingen

Welke moeilijkheidsgraad wil je voor de volgende oefeningenreeks?

Heel makkelijk, Makkelijk, Gewoon, Moeilijk, Heel moeilijk

ingen in de reeks juist oplost, dan stijgt je level:

Expert, Bedreven, Competent, Gevorderde beginner, Beginner

Start de reeks

aangeraden oefening van hetzelfde hoofdstuk

Aangeraden: Oefening 37, Oefening 26, Oefening 21

Hoe is je nieuw niveau bepaald?

Wiski schat jouw niveau en de moeilijkheid van oefeningen in. Beide veranderen bij het je niveau is gestegen na het maken van de reeks oefeningen. Daarna is het nog extra gestegen door je feedback.

Maak oefening 21

... of kies zelf je volgende oefening

Naar het oefeningenoverzicht



How good do you think you are at math?

Expert: mathematics holds no secrets for you

Proficient: you score better than average on math

Competent: you score average on math

Advanced beginner: basic exercises are no

Novice: you often have a hard time understand

Submit

Patrol ID: 3033, Age: 47, Blood Sugar: 7.8 g, Drinking Status: Rarely Drinks, Physical Activity Level: Low

Patrol Summary: Blood Sugar: 7.5, Waist Masser: 112, BMI: 33.9

Recommendations to reduce risk: Exercise everyday for 30 min, Before walk increase by 14 cm

Impactful Risk Factors: Blood Sugar, Waist Masser, BMI

Risk Recovery: Blood Sugar, Waist Masser, BMI