

**ARTICLE: COUNTERFACTUAL EXPLANATIONS WITHOUT OPENING THE  
BLACK BOX: AUTOMATED DECISIONS AND THE GDPR**

Spring, 2018

**Reporter**

31 Harv. J. Law & Tec 841 \*

**Length:** 11071 words

**Author:** Sandra Wachter, \* Brent Mittelstadt, \*\* & Chris Russell \*\*\*

\* Oxford Internet Institute, University of Oxford, 1 St. Giles, Oxford, OX1 3JS, UK and The Alan Turing Institute, British Library, 96 Euston Road, London, NW1 2DB, UK. Email: [sandra.wachter@oii.ox.ac.uk](mailto:sandra.wachter@oii.ox.ac.uk). This work was supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1.

\*\* Oxford Internet Institute, University of Oxford, 1 St. Giles, Oxford, OX1 3JS, UK, The Alan Turing Institute, British Library, 96 Euston Road, London, NW1 2DB, UK, Department of Science and Technology Studies, University College London, 22 Gordon Square, London, WC1E 6BT, UK.

\*\*\* The Alan Turing Institute, British Library, 96 Euston Road, London, NW1 2DB, UK, Department of Electrical and Electronic Engineering, University of Surrey, Guildford, GU2 7HX, UK.

The authors would like to thank Patrick Allo, Dan Burk, Lee Bygrave, Mireille Hildebrandt, David Landsberg, Zachary Lipton, Alessandro Spina, David Sutcliffe, Jade Thompson, Jessie Wand, Tal Zarsky, participants of the 2018 Privacy Law Scholars Conference (Brussels), the 2017 Ethical Auditing in Automated Decision-Making workshop (Oxford), and the 2018 Artificial Intelligence, Ethics and the Law workshop (London), and members of the Alan Turing Institute's Fairness, Transparency and Privacy interest group for their invaluable feedback and comments. We would also like to thank the Oxford Internet Institute and Alan Turing Institute for providing the highly interdisciplinary environments in which the authors were able to develop the central idea of the paper. Last but not least we want to thank the editorial team of Harvard Journal of Law and Technology for their invaluable feedback and insightful comments that massively improved the quality of our work.

**Text**

---

[\*842] I. INTRODUCTION

There has been much discussion of the existence of a "right to explanation" in the EU General Data Protection Regulation ("GDPR"), and its merits and disadvantages.<sup>1</sup> Attempts to implement a right to explanation that opens the "black box" to provide insight into the internal decision-making process of algorithms face four major legal and technical barriers. First, a legally binding right to explanation does not exist in the GDPR.<sup>2</sup> Second, even if legally binding, the right would only apply in limited cases (when a negative decision was solely automated and had legal or other similar significant effects).<sup>3</sup> Third, explaining the functionality of complex algorithmic decision-making systems and their rationale in specific cases is a technically challenging problem.<sup>4</sup> Explanations may likewise offer [\*843] little meaningful information to data subjects, raising questions about their value.<sup>5</sup> Finally, data controllers have an interest in not sharing details of their algorithms to avoid disclosing trade secrets, violating the rights and freedoms of others (e.g. privacy), and allowing data subjects to game or manipulate the decision-making system.<sup>6</sup>

Despite these difficulties, the social and ethical value (and perhaps responsibility) of offering explanations to affected data subjects remains unaffected. One significant point has been neglected in this discussion. An explanation of automated decisions, both as envisioned by the GDPR and in general, does not necessarily hinge on the general public understanding of how algorithmic systems function. Even though such interpretability is of great importance and should be pursued, explanations can, in principle, be offered without opening the "black box." Looking at explanations as a means to help a data subject *act* rather than merely understand, one could gauge the scope and content of explanations according to the specific goal or action they are intended to support.

---

<sup>1</sup> See, e.g., Sandra Wachter, Brent Mittelstadt & Luciano Floridi, *Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation*, 7 INT'L DATA PRIVACY L. 76, 79-90 (2017); Isak Mendoza & Lee A. Bygrave, *The Right Not to Be Subject to Automated Decisions Based on Profiling*, in EU INTERNET LAW: REGULATION AND ENFORCEMENT (Tatiani Synodinou et al. eds., 2017), <https://papers.ssrn.com/abstract=2964855> [<https://perma.cc/XV3T-G98W>]; Lilian Edwards & Michael Veale, *Slave to the Algorithm? Why a 'Right to Explanation' is Probably Not the Remedy You are Looking For*, 16 DUKE L. TECH. REV. 18, 18-19 (2017); Gianclaudio Malgieri & Giovanni Comand , *Why a Right to Legibility of Automated Decision-Making Exists in the General Data Protection Regulation*, 7 INT'L DATA PRIVACY L. 243, 246-47 (2017); Andrew Selbst & Julia Powles, *Meaningful Information and the Right to Explanation*, 7 INT'L DATA PRIV. L. 233, 233-34 (2017); Bryce Goodman & Seth Flaxman, *EU Regulations on Algorithmic Decision-Making and a "Right to Explanation,"* ARXIV, Aug. 31, 2016, at 6-7, <http://arxiv.org/abs/1606.08813> [<https://perma.cc/5ZTR-WG8R>]; Tae Wan Kim & Bryan Routledge, *Algorithmic Transparency, a Right to Explanation, and Placing Trust*, SQUARESPACE (June 2017), <https://static1.squarespace.com/static/592ee286d482e908d35b8494/t/59552415579fb30c014cd06c/1498752022120/Algorithmic+transparen+cy%2C+a+right+to+explanation+and+trust+%28TWK%26BR%29.pdf> [<https://perma.cc/K53W-GVN2>].

<sup>2</sup> See Kim & Routledge, *supra* note 1, at 3; Wachter, Mittelstadt & Floridi, *supra* note 1, at 79.

<sup>3</sup> See Wachter, Mittelstadt & Floridi, *supra* note 1, at 78.

<sup>4</sup> See, e.g., Wachter, Mittelstadt & Floridi, *supra* note 1, at 77; Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633, 638 (2016); Tal Zarsky, *Transparent Predictions*, 2013 U. ILL. L. REV. 1503, 1519-20 (2013); Edwards & Veale, *supra* note 1, at 22.

<sup>5</sup> See Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, 3 BIG DATA & SOC., Jan. 5, 2016, at 5; Kroll et al., *supra* note 4, at 638.

<sup>6</sup> See Roger A. Ford & W. Nicholson Price II, *Privacy and Accountability in Black-Box Medicine*, 23 MICH. TELECOMM. TECH. REV. 1, 3 (2016); Mike Ananny & Kate Crawford, *Seeing Without Knowing: Limitations of the Transparency Ideal and its Application to Algorithmic Accountability*, 20 NEW MEDIA & SOC., Dec. 13, 2016, at 8, <http://journals.sagepub.com/doi/full/10.1177/1461444816676645> [<https://perma.cc/3HF6-G9DS>]; Burrell, *supra* note 5, at 3; Frank A. Pasquale, *Restoring Transparency to Automated Authority*, 9 J. TELECOMM. HIGH TECH. L. 235, 237 (2011); Brenda Reddix-Small, *Credit Scoring and Trade Secrecy: An Algorithmic Quagmire or How the Lack of Transparency in Complex Financial Models Scuttled the Finance Market*, 12 U.C. DAVIS BUS. L.J. 87, 94 (2011).

Explanations can serve many purposes. To investigate the potential scope of explanations, it seems reasonable to start from the perspective of the data subject, which is the natural person whose data is being collected and evaluated. We propose three aims for explanations to assist data subjects: (1) to inform and help the subject understand why a particular decision was reached, (2) to provide grounds to contest adverse decisions, and (3) to understand what could be changed to receive a desired result in the future, based on the current decision-making model. As we show, the GDPR offers little support to achieve any of these aims. However, none hinge on explaining the internal logic of automated decision-making systems.

Building trust is essential to increase societal acceptance of algorithmic decision-making. To close current gaps in transparency and accountability that undermine trust between data controllers and data subjects,<sup>7</sup> we propose to move beyond the limitations of the GDPR. [\*844] We argue that counterfactuals should be used as a means to provide explanations for individual decisions.

*Unconditional counterfactual explanations* should be given for positive and negative automated decisions, regardless of whether the decisions are solely (as opposed to predominantly) automated or produce legal or other significant effects. This approach provides data subjects with meaningful explanations to understand a given decision, grounds to contest it, and advice on how the data subject can change his or her behaviour or situation to possibly receive a desired decision (e.g. loan approval) in the future without facing the severely limited applicability imposed by the GDPR's definition of automated individual decision-making.<sup>8</sup>

In this paper, we present the concept of unconditional counterfactual explanations as a novel type of explanation of automated decisions that overcomes many challenges associated with algorithmic interpretability and accountability. We situate counterfactuals in the philosophical history of knowledge, as well as historical and modern research on interpretability and fairness in machine learning. Based on the potential advantages offered to data subjects by counterfactual explanations, we then assess their alignment with the GDPR's numerous provisions concerning automated decision-making. Specifically, we examine whether the GDPR offers support for explanations that aim to help data subjects understand the scope of automated decisionmaking as well as the rationale of specific decisions, explanations to contest decisions, and explanations that offer guidance on how data subjects can change their behaviour to receive a desired result. We conclude that unconditional counterfactual explanations can bridge the gap between the interests of data subjects and data controllers that otherwise acts as a barrier to a legally binding right to explanation.

## II. COUNTERFACTUALS

Counterfactual explanations take a similar form to the statement:

You were denied a loan because your annual income was £ 30,000. If your income had been £ 45,000, you would have been offered a loan.

Here the statement of decision is followed by a counterfactual, or statement of how the world would have to be different for a desirable outcome to occur. Multiple counterfactuals are possible, as multiple desirable outcomes can exist, and there may be several ways to [\*845] achieve any of these outcomes. The concept of the "closest possible world," or the smallest change to the world that can be made to obtain a desirable

---

<sup>7</sup> See Mendoza & Bygrave, *supra* note 1, at 97; Wachter, Mittelstadt & Floridi, *supra* note 1, at 78.

<sup>8</sup> See Wachter, Mittelstadt & Floridi, *supra* note 1, at 87-88; Edwards & Veale, *supra* note 1, at 22; Mendoza & Bygrave, *supra* note 1, at 83.

outcome, is key throughout the discussion of counterfactuals. In many situations, providing several explanations covering a range of diverse counterfactuals corresponding to relevant or informative "close possible worlds" rather than "the closest possible world" may be more helpful. At the same time, knowing the smallest possible change to a variable or set of variables to arrive at a different outcome may not always be the most helpful type of counterfactual. Rather, relevance will depend also upon other case-specific factors, such as the mutability of a variable or real world probability of a change.<sup>9</sup>

In the existing literature, "explanation" typically refers to an attempt to convey the internal state or logic of an algorithm that leads to a decision.<sup>10</sup> In contrast, counterfactuals describe a dependency on the external facts that led to that decision. This is a crucial distinction. In modern machine learning, the internal state of the algorithm can consist of millions of variables intricately connected in a large web of dependent behaviours.<sup>11</sup> Conveying this state to a layperson in a way that allows them to reason about the behaviour of an algorithm is extremely challenging.<sup>12</sup>

The machine learning and legal communities have both taken relatively restricted views on what passes for an explanation. The machine learning community has been primarily concerned with debugging<sup>13</sup> and conveying approximations of algorithms that programmers or researchers could use to understand which features are important<sup>14</sup> while law and ethics scholars have been more concerned with understanding the internal logic of decisions as a means to assess [\*846] their lawfulness (e.g. prevent discriminatory outcomes), contest them, increase accountability generally, and clarify liability.<sup>15</sup>

---

<sup>9</sup> See *infra* Section II.A.

<sup>10</sup> See Burrell, *supra* note 5, at 1.

<sup>11</sup> See, e.g., Kaiming He et al., *Deep Residual Learning for Image Recognition*, in PROCEEDINGS OF THE IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION 770-78 (2016).

<sup>12</sup> See Burrell, *supra* note 5, at 1; Zachary C. Lipton, *The Mythos of Model Interpretability*, in 2016 WORKSHOP ON HUMAN INTERPRETABILITY IN MACHINE LEARNING 96, [https://zacklipton.com/media/papers/mythos\\_model\\_interpretability\\_lipton2016.pdf](https://zacklipton.com/media/papers/mythos_model_interpretability_lipton2016.pdf) [<https://perma.cc/4JVZ-7T6D>].

<sup>13</sup> Osbert Bastani, Carolyn Kim & Hamsa Bastani, *Interpretability via Model Extraction*, ARXIV:1706.09773, Mar. 13, 2018, at 1, <https://arxiv.org/pdf/1706.09773.pdf> [<https://perma.cc/3GAJ-8GE9>].

<sup>14</sup> Marco Tulio Ribeiro, Sameer Singh & Carlos Guestrin, *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 1135 (2016); Ramprasaath R. Selvaraju et al., *Grad-CAM: Why Did You Say That?*, ARXIV, Nov. 22, 2016, at 1, <https://arxiv.org/abs/1611.07450> [<https://perma.cc/AA8F-45XJ>]; Karen Simonyan, Andrea Vedaldi & Andrew Zisserman, *Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps*, ARXIV, Dec. 20, 2013, at 1, <https://arxiv.org/abs/1312.6034> [<https://perma.cc/Y85R-X9UE>].

<sup>15</sup> See, e.g., Finale Doshi-Velez et al., *Accountability of AI Under the Law: The Role of Explanation*, ARXIV, Nov. 3, 2017, at 1, <https://arxiv.org/pdf/1711.01134.pdf> [<https://perma.cc/JGN7-78WR>]; Finale Doshi-Velez, Ryan Budish & Mason Kortz, *The Role of Explanation in Algorithmic Trust*, TRUSTWORTHY ALGORITHMIC DECISION-MAKING 2, <http://trustworthy-algorithms.org/whitepapers/Finale%20Doshi-Velez.pdf> [<https://perma.cc/4L88-V58A>]; Mireille Hildebrandt, *The Dawn of a Critical Transparency Right for the Profiling Era*, in DIGITAL ENLIGHTENMENT YEARBOOK 2012 41 (Jacques Bus et al. eds., 2012); Tim Miller, *Explanation in Artificial Intelligence: Insights from the Social Sciences*, ARXIV, Jun. 22, 2017, at 3, <https://arxiv.org/pdf/1706.07269.pdf> [<https://perma.cc/WF2G-4NUS>]; Pasquale, *supra* note 6, at 236; Danielle Keats Citron & Frank A. Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 6-7 (2014), <https://papers.ssrn.com/abstract=2376209> [<https://perma.cc/9CXY-DBTN>]; Tal Zarsky, *Transparent Predictions*, 2013 U. ILL. L. REV. 1503, 1506-09 (2013), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2324240](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2324240)

As such, the proposal made here for counterfactuals as explanations lies outside of the taxonomies of explanations proposed previously in machine learning, legal, and ethical literature. In contrast, as we discuss in the next Section, analytic philosophy has taken a much broader view of knowledge and how counterfactuals can be used as justifications of beliefs.<sup>16</sup>

### A. Historic Context and the Problem of Knowledge

Analytic philosophy has a long history of analysing the necessary conditions for propositional knowledge.<sup>17</sup> Expressions of the type "*S* knows that *p*" constitute knowledge, where *S* refers to the knowing subject, and *p* to the proposition that is known. Traditional approaches, which conceive of knowledge as "justified true belief," conceive of three necessary conditions for knowledge: truth, belief, and justification.<sup>18</sup> According to this tripartite approach, in order to know something, it is not enough to simply believe that something is true: rather, you must also have a good reason for believing it.<sup>19</sup> The relevance of this approach comes from the observation that this form of justification of beliefs can serve as a type of explanation,<sup>20</sup> as it is fundamentally a reason that a belief is held and therefore serves as an answer to [\*847] the question, "Why do you believe *X*?" Understanding the different forms these justifications can take opens the door to a broader class of explanations than previously encountered in interpretability research.

Although influential, "justified true belief" has faced much criticism<sup>21</sup> and inspired substantial analysis of modifications to this tripartite approach as well as proposals for additional necessary conditions for a proposition to constitute knowledge.<sup>22</sup> Modal conditions, including safety<sup>23</sup> and sensitivity,<sup>24</sup> have been proposed as necessary additions to the tripartite built on counterfactual relations.<sup>25</sup>

---

[<https://perma.cc/F8FC-YDJG>]; Tal Zarsky, *The Trouble with Algorithmic Decisions: An Analytic Road Map to Examine Efficiency and Fairness in Automated and Opaque Decision Making*, 41 SCI. TECH. HUM. VALUES 118, 118-32 (2016).

<sup>16</sup> See, e.g., DAVID LEWIS, COUNTERFACTUALS 1-4, 84-91 (1973); ROBERT NOZICK, PHILOSOPHICAL EXPLANATIONS 172-74 (1981); David Lewis, *Counterfactuals and Comparative Possibility*, 2 J. PHIL. LOGIC 418, 418-46 (1973); Peter Lipton, *Contrastive Explanation*, 27 ROYAL INST. PHIL. SUPP. 247, 247 (1990).

<sup>17</sup> See generally ALFRED JULES AYER, THE PROBLEM OF KNOWLEDGE (1956).

<sup>18</sup> See generally Edmund L. Gettier, *Is Justified True Belief Knowledge?*, 23 ANALYSIS 6, 121 (1963); Julien Dutant, *The Legend of the Justified True Belief Analysis*, 29 PHIL. PERSP. 1, 95 (2015).

<sup>19</sup> See Gettier, *supra* note 18, at 121.

<sup>20</sup> See NOZICK, *supra* note 16, at 174.

<sup>21</sup> See, e.g., Dutant, *supra* note 18, at 95; Mark Kaplan, *It's Not What You Know That Counts*, 82 J. PHIL. 350, 350 (1985).

<sup>22</sup> Jonathan Ichikawa & Matthias Steup, *The Analysis of Knowledge*, in STANFORD ENCYCLOPEDIA OF PHILOSOPHY ARCHIVE (Fall 2017 ed.), <https://plato.stanford.edu/archives/fall2017/entries/knowledge-analysis> (last visited May 6, 2018).

<sup>23</sup> See Ernest Sosa, *How to Defeat Opposition to Moore*, 13 PHIL. PERSP. 141, 141-43 (1999).

<sup>24</sup> See Jonathan Ichikawa, *Quantifiers, Knowledge, and Counterfactuals*, 82 PHIL. PHENOMENOLOGICAL RES. 287, 287 (2011); see also NOZICK, *supra* note 16, at 172-74.

<sup>25</sup> See Ichikawa & Steup, *supra* note 22, Section 5 (reviewing these concepts and their criticisms).

Sosa <sup>26</sup> as well as Ichikawa and Steup <sup>27</sup> define sensitivity as:

If  $p$  were false,  $S$  would not believe  $p$ .

Here, the statement "If  $p$  were false" is a counterfactual defining a "possible world" close to the world in which  $p$  is true. <sup>28</sup> The sensitivity condition suggests that "in the nearest possible worlds in which not- $p$ , the subject does not believe that  $p$ ." <sup>29</sup> Our notion of counterfactual explanations hinges upon the related concept:

If  $q$  were false,  $S$  would not believe  $p$ .

We claim that in this case,  $q$  serves as an explanation of  $S$ 's belief in  $p$ , inasmuch as  $S$  only holds belief  $p$  while  $q$  is true, and that changing  $q$  would also cause  $S$ 's belief to change. A key point is that such statements only describe  $S$ 's beliefs, which need not reflect reality. <sup>30</sup> As such, these statements can be made without knowledge of any causal relationship between  $q$  and  $p$ .

We define counterfactual explanations as statements taking the form:

**[\*848]** Score  $p$  was returned because variables  $V$  had values ( $v[1]$ ,  $v[2]$ , . . .) associated with them. If  $V$  instead had values ( $v[1]'$ ,  $v[2]'$ , . . .), and all other variables had remained constant, score  $p'$  would have been returned.

While many such explanations are possible, an ideal counterfactual explanation would alter values as little as possible and represent a closest world under which score  $p'$  is returned instead of  $p$ . The notion of a "closest possible world" is thus implicit in our definition.

Our version of counterfactuals perhaps most resembles a structural equations approach in execution by identifying alterations to variables. This approach is more similar to Pearl's "mini-surgeries" <sup>31</sup> than Lewis' "miracles." <sup>32</sup> In any case, our approach does not rely on knowledge of the causal structure of the world, <sup>33</sup> or suggest which context-dependent metric of distance between worlds is preferable to establish causality. <sup>34</sup> In many situations, it will be more informative to provide a diverse set of counterfactual explanations, corresponding to different choices of nearby possible worlds for which the counterfactual holds or a preferred outcome is delivered, rather than a theoretically ideal counterfactual describing the "closest

---

<sup>26</sup> Sosa, *supra* note 23, at 141.

<sup>27</sup> Ichikawa & Steup, *supra* note 22, Section 5.1.

<sup>28</sup> See LEWIS, *supra* note 16, at 1-4.

<sup>29</sup> Ichikawa & Steup, *supra* note 22, Section 5.1.

<sup>30</sup> For example,  $S$  could believe that a person is inherently more trustworthy ( $p$ ) because they are a Capricorn ( $q$ ).

<sup>31</sup> See JUDEA PEARL, CAUSATION 223-24 (2000).

<sup>32</sup> LEWIS, *supra* note 16, at 47-48.

<sup>33</sup> See *infra* Section II.D.

<sup>34</sup> See Boris Kment, *Counterfactuals and Explanation*, 115 MIND 261, 261-309 (2006).

possible world" according to a preferred distance metric.<sup>35</sup> Case-specific considerations will be relevant to the choice of distance metric and a "sufficient" and "relevant" set of counterfactual explanations. Such considerations may include the capabilities of the individual concerned, sensitivity, mutability of the variables involved in a decision, and ethical or legal requirements for disclosure.<sup>36</sup>

Similarly, counterfactuals that describe changes to multiple variables within the model can be provided. These would represent possible futures brought about by changes to the individual's circumstances. As an example, the impact of changes in income could be calculated in combination with changes to career, thereby ensuring the counterfactual represents a realistic possible world.

**[\*849]** *B. Explanations in A.I. and Machine Learning*

Much of the early work in A.I. on explaining the decisions made by expert or rule-based systems focused on classes of explanation closely related to counterfactuals. For example, Gregor and Benbasat<sup>37</sup> offer the following example of what they call a type 1 explanation:

Q: Why is a tax cut appropriate?

A: Because a tax cut's preconditions are high inflation and trade deficits, and current conditions exhibit these factors.

Buchanan and Shortliffe<sup>38</sup> offer a similar example:

RULE009 IF:

- 1) The gram stain of the organism is gramneg, and
- 2) The morphology of the organism is coccus

THEN: There is strongly suggestive evidence (.8) that the identity of the organism is Neisseria

As is typical in early A.I., questions we now recognise as hard such as "How do we decide if inflation is high?" or "Why are these the preconditions of a tax cut?" are assumed to have been addressed by humans, and are not discussed as part of the explanation.<sup>39</sup> As such, the explanations do not provide insight into what people in machine learning think of as the internal logic of black box classifiers. Instead, the first example can be rewritten as two diverse counterfactual statements:

If inflation was lower, a tax cut would not be recommended.

---

<sup>35</sup> The merits of different metrics of distance between possible worlds have long been debated in philosophy without the emergence of consensus. Meaningfully addressing this debate goes beyond the scope of this paper which proposes a method for counterfactual explanations, but will be explored in future work. For further discussion of distance metrics and counterfactuals, see LEWIS, *supra* note 16, at 8-15; Ernest W. Adams, *On the Rightness of Certain Counterfactuals*, 74 PAC. PHIL. Q. 1, 1-8 (1993); Kment, *supra* note 34, at 262.

<sup>36</sup> A discussion of appropriate metrics for making these choices goes beyond the scope of this paper, but will be addressed in future work. With that said, relevant philosophical discussion can be found on determining relevance of possible causal or contrastive explanations, counterfactuals, and distance metrics. See, e.g., Adams, *supra* note 35, at 1-8; Peter Lipton, *Contrastive Explanation*, 27 ROYAL INST. PHIL. SUPP. 247, 254-65 (1990).

<sup>37</sup> Shirley Gregor & Izak Benbasat, *Explanations from Intelligent Systems: Theoretical Foundations and Implications for Practice*, 23 MIS Q. 497, 503 (1999).

<sup>38</sup> BRUCE G. BUCHANAN & EDWARD D. SHORTLIFFE, *RULE-BASED EXPERT SYSTEMS: THE MYCIN EXPERIMENTS OF THE STANFORD HEURISTIC PROGRAMMING PROJECT* 344 (1984).

<sup>39</sup> See, e.g., Gregor & Benbasat, *supra* note 37, at 503.

If there was no trade deficit, a tax cut would not be recommended.

[\*850] While the second example is closely related to the counterfactual <sup>40</sup>:

If the gram stain was negative or the morphology was not coccus, the algorithm would not be confident that the organism is *Neisseria*.

The most important difference between these earlier approaches and counterfactuals is that counterfactuals continue functioning in an end-to-end integrated approach. If the gram stain and morphology in the MYCIN example were also determined by the algorithm, counterfactuals would automatically return a close sample with a different classification, while these early methods could not be applied to such involved scenarios.

As focus has switched from A.I. and logic-based systems towards machine learning tasks such as image recognition, the notion of an explanation has come to refer to providing insight into the internal state of an algorithm, or to human-understandable approximations of the algorithm. <sup>41</sup> As such, the most related machine learning work to these, and to ours, is by Martens and Provost. <sup>42</sup> Uniquely among other works in machine learning, their work shares our interest in making interventions to alter the outcome of classifier responses. However, the work is firmly linked to the problem of document classification, and the only interventions it proposes involve the removal of words from documents to stop websites from being classified as "adult." <sup>43</sup> The heuristic proposed cannot be easily generalised to either continuous variables, <sup>44</sup> or even the addition of words to documents.

The majority of works in machine learning on explanations and interpreting models concern themselves with generating simple models as local approximations of decisions. <sup>45</sup> Generally, the idea is to create a simple human-understandable approximation of a decisionmaking algorithm that accurately models the decision given the current [\*851] inputs, but may be arbitrarily bad for different inputs. <sup>46</sup> However, there are numerous difficulties with treating these approaches as explanations suitable for a lay data subject.

In general, it is unclear if these models are interpretable by nonexperts. They make a three-way trade-off between the quality of the approximation, the ease of understanding the function, and the size of the domain for which the approximation is valid. <sup>47</sup> As we show in Appendix 1, these local models can produce widely

---

<sup>40</sup> However, they are not logically equivalent. The example from MYCIN differs in that it is still possible that some samples that are either gram positive or have a different morphology could still be classified as *Neisseria*.

<sup>41</sup> See Ribeiro et al., *supra* note 14, at 1135-37.

<sup>42</sup> See David Martens & Foster Provost, *Explaining Data-Driven Document Classifications*, 38 MIS Q. 73, 73-74 (2013).

<sup>43</sup> *Id.*

<sup>44</sup> "Continuous variables" refers to variables whose assigned values are not restricted to a small set of discrete values: such as 'present' or 'not present', but instead can take any value in a given range. Measurements such as height, weight, or how bright a particular pixel is in a photo, are often treated as continuous variables. See, e.g., *Maths: Discrete and Continuous Variables*, BBC (2014), <http://www.bbc.co.uk/schools/gcsebitesize/maths/statistics/samplinghirev1.shtml> [<https://perma.cc/8B7M-NJ39>].

<sup>45</sup> See Ribeiro et al., *supra* note 14, at 1135; Selvaraju et al., *supra* note 14, at 1-3; Simonyan et al., *supra* note 14, at 1.

<sup>46</sup> See Ribeiro et al., *supra* note 14, at 1143; Selvaraju et al., *supra* note 14, at 1-3; Simonyan et al., *supra* note 14, at 1.

<sup>47</sup> See Bastani et al., *supra* note 13, at 1; Himabindu Lakkaraju et al., *Interpretable & Explorable Approximations of Black Box Models*, ARXIV, July 4, 2017, at 1, <https://arxiv.org/pdf/1707.01154.pdf> [<https://perma.cc/6JFE-N4YD>].



varying estimates of the importance of variables even in simple scenarios such as the single variable case, making it extremely difficult to reason about how a function varies as the inputs change. Moreover, the utility of such approaches outside of model debugging by expert programmers is unclear. Research has yet to be conducted on how to convey the various limitations and unreliabilities of these approaches to a lay audience in such a way that they can make use of such explanations.

In contrast, counterfactual explanations are intentionally restricted. They are crafted in such a way as to provide a minimal amount of information capable of altering a decision, and they do not require the data subject to understand any of the internal logic of a model in order to make use of it. The downside to this is that individual counterfactuals may be overly restrictive. A single counterfactual may show how a decision is based on certain data that is both correct and unable to be altered by the data subject before future decisions, even if other data exist that could be amended for a favourable outcome. This problem could be resolved by offering multiple diverse counterfactual explanations to the data subject.

### *C. Adversarial Perturbations and Counterfactual Explanations*

The techniques used to generate counterfactual explanations on deep networks such as resnet<sup>48</sup> are already widely studied in the machine learning literature under the name "Adversarial Perturbations."<sup>49</sup> In these works, algorithms capable of computing [\*852] counterfactuals are used to confuse existing classifiers by generating a synthetic data point close to an existing one such that the new synthetic data point is classified differently than the original one.<sup>50</sup>

One strength of counterfactuals is that they can be efficiently and effectively computed by applying standard techniques, even to cutting-edge architectures. Some of the largest and deepest neural networks are used in the field of computer vision, particularly in image labelling tasks such as ImageNet.<sup>51</sup> These types of classifiers have been shown to be particularly vulnerable to Adversarial Perturbation attacks, where small changes to a given image can result in the image being assigned to an entirely different class. For example, DeepFool<sup>52</sup> defines an adverse perturbation of an image  $x$ , given a classifier, as the smallest change to  $x$  such that the classification changes. Essentially, this is a counterfactual by a different name. Finding a closest possible world to  $x$  such that the classification changes is, under the right choice of distance function, the same as finding the smallest change to  $x$ .

---

<sup>48</sup> See He et al., *supra* note 11, at 770.

<sup>49</sup> See Ian J. Goodfellow, Jonathon Shlens & Christian Szegedy, *Explaining and Harnessing Adversarial Examples*, ArXiv, Mar. 20, 2015, at 1, <https://arxiv.org/pdf/1412.6572.pdf> [<https://perma.cc/64BR-WVE7>]; Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi & Pascal Frossard, *Deepfool: A Simple and Accurate Method to Fool Deep Neural Networks*, in PROCEEDINGS OF THE IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION 2574-82 (2016); Christian Szegedy et al., *Intriguing Properties of Neural Networks*, ARXIV, Feb. 19, 2014, at 2, <https://arxiv.org/pdf/1312.6199.pdf> [<https://perma.cc/K37R-6NP2>].

<sup>50</sup> See Goodfellow, Shlens & Szegedy, *supra* note 49, at 1; Moosavi-Dezfooli, Fawzi & Frossard, *supra* note 49, at 2574-82; Szegedy et al., *supra* note 49, at 2.

<sup>51</sup> See Jia Deng et al., *Imagenet: A Large-Scale Hierarchical Image Database*, in IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION 248-55 (2009).

<sup>52</sup> See Moosavi-Dezfooli, Fawzi & Frossard, *supra* note 49, at 2574.

Importantly, none of the standard works on Adversarial Perturbations make use of appropriate distance functions, and the majority of such approaches tend to favour making small changes to many variables, instead of providing sparse human interpretable solutions that modify only a few variables.<sup>53</sup> Despite this, efficient computation of counterfactuals and Adversarial Perturbations is made possible by virtue of state-of-the-art algorithms being differentiable. Many optimisation techniques proposed in the Adversarial Perturbation literature are directly applicable to this problem, making counterfactual generation efficient.

One of the more challenging aspects of Adversarial Perturbations is that these small perturbations of an image are barely human-perceptible, but result in drastically different classifier responses.<sup>54</sup> Informally, this appears to happen because the newly generated images do not lie in the "space of real-images," but slightly outside it.<sup>55</sup> This phenomenon serves as an important reminder that when computing counterfactuals by searching for a close possible world, it is at least as important that the solution found comes from a "possible world" as it is that it is close to the starting example. Further research into how data from high-dimensional and highly-structured spaces, such as natural images, can be characterised is needed before counterfactuals can be reliably used as explanations in these spaces.

#### *D. Causality and Fairness*

Several works have approached the problem of guaranteeing that algorithms are fair, i.e. that they do not exhibit a bias towards particular ethnic, gender, or other protected groups, using causal reasoning<sup>56</sup> and counterfactuals.<sup>57</sup> Kusner et al.<sup>58</sup> consider counterfactuals where the subject belongs to a different race or sex, and require that the decision made remain the same under such a counterfactual for it to be considered fair. In contrast, we consider counterfactuals in which the decision differs from its current state.

Many works have suggested that transparency might be a useful tool for enforcing fairness. While it is unclear how counterfactuals could be used for this purpose, it is also unclear if any form of explanation of individual decisions can in fact help. Grgic-Hlaca et al.<sup>59</sup> showed how understandable models can easily mislead our intuitions, and that predominantly using features people believed to be fair slightly *increased* the racism exhibited by algorithms, while decreasing accuracy. In general, the best tools for uncovering

---

<sup>53</sup> See Jiawei Su, Danilo Vasconcellos Vargas & Sakurai Kouichi, *One Pixel Attack for Fooling Deep Neural Networks*, ARXIV, Nov. 16, 2017, at 8-9, <http://arxiv.org/abs/1710.08864> [<https://perma.cc/5F2N-JBJF>].

<sup>54</sup> Niki Kilbertus et al., *Avoiding Discrimination through Causal Reasoning*, ARXIV, Jan. 21, 2018, at 1, <https://arxiv.org/pdf/1706.02744.pdf> [<https://perma.cc/8GZQ-PRSJ>].

<sup>55</sup> See generally Simant Dube, *High Dimensional Spaces, Deep Learning and Adversarial Examples*, ARXIV, Jan. 14, 2018, <https://arxiv.org/pdf/1801.00634.pdf> [<https://perma.cc/8FE5-RY4X>] (presenting preliminary investigation of this matter).

<sup>56</sup> See Kilbertus et al., *supra* note 54, at 1.

<sup>57</sup> See Matt J. Kusner et al., *Counterfactual Fairness*, ARXIV, Mar. 8, 2018, at 16, <https://arxiv.org/pdf/1703.06856.pdf> [<https://perma.cc/4SYN-7J9D>].

<sup>58</sup> *Id.*

<sup>59</sup> Nina Grgic-Hlaca et al., *The Case for Process Fairness in Learning: Feature Selection for Fair Decision Making*, in NIPS SYMPOSIUM ON MACHINE LEARNING AND THE LAW 8 (2016).

systematic biases are likely to be based upon large-scale statistical analysis and not upon explanations of individual decisions.<sup>60</sup>

With that said, counterfactuals can provide evidence that an algorithmic decision is affected by a protected variable (e.g. race), and that it may therefore be discriminatory.<sup>61</sup> For the types of distance function we consider in the next Section, if the counterfactuals found change someone's race, then the treatment of that individual is dependent on race. However, the converse statement is *not* true. Counterfactuals that do not modify a protected attribute cannot be used as evidence [\*854] that the attribute was irrelevant to the decision. This is because counterfactuals describe only some of the dependencies between a particular decision and specific external facts. This can be seen clearly in Section III.A, where the counterfactuals proposed for a particular classifier involve 'black' people changing their race, while not suggesting that 'white' people's race should be varied.

### III. GENERATING COUNTERFACTUALS

In the following Section, we give examples of how meaningful counterfactuals can be easily computed. Many of the standard classifiers of machine learning (including Neural Networks, Support Vector Machines, and Regressors) are trained by finding the optimal set of weights  $w$  that minimises an objective over a set of training data:

Equation 1

Equation 1

Where  $y_i$  is the label for data point  $x_i$  and  $p(\cdot)$  is a regularizer over the weights. We wish to find a counterfactual  $x'$  as close to the original point  $x[i]$  as possible such that  $f_w(x')$  is equal to a new target  $y'$ . We can find  $x'$  by holding  $w$  fixed and minimizing the related objective:

Equation 2

Where  $d(\cdot, \cdot)$  is a distance function that measures how far the counterfactual  $x'$  and the original data point  $x_i$  are from one another. In practice, maximisation over  $\lambda$  is done by iteratively solving for  $x'$  and increasing  $\lambda$  until a sufficiently close solution is found.

The choice of optimiser for these problems is relatively unimportant. In practice, any optimiser capable of training the classifier under Equation 1 seems to work equally well, and we use ADAM<sup>62</sup> for all experiments. As local minima are a concern, we initialise each run with different random values for  $x'$  and select as our counterfactual the best minimizer of Equation 2. These different minima can be used as a diverse set of multiple counterfactuals.

---

<sup>60</sup> See Andrea Romei & Salvatore Ruggieri, *A Multidisciplinary Survey on Discrimination Analysis*, 29 KNOWLEDGE ENGINEERING REV. 582, 617 (2014).

<sup>61</sup> Establishing the influence of a protected variable on a decision does not, by itself, prove that illegal discrimination has occurred. Mitigating factors may exist which justify the usage of a protected attribute. See, e.g., Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CAL. L. REV. 671, 676 (2016) (discussing disparate treatment in American anti-discrimination law).

<sup>62</sup> Diederik Kingma & Jimmy Ba, *ADAM: A Method for Stochastic Optimization*, ARXIV, Dec. 22, 2014, at 1-4, <https://arxiv.org/pdf/1412.6980.pdf> [<https://perma.cc/3RH4-WSXG>].

[\*855] Of particular importance is the choice of distance function used to decide which synthetic data point  $x'$  is closest to the original data point  $x[i]$ . As a sensible first choice, which should be refined based on subject- and task-specific requirements, we suggest use of the  $L[1]$  norm, or Manhattan distance, weighted by the inverse median absolute deviation. This is written as  $MAD[k]$  for the median absolute deviation of feature  $k$ , over the set of points  $P$ :

Equation 3

We chose  $d(.,.)$  as:

Equation 4

This distance metric has several desirable properties. Firstly, it captures some of the intrinsic volatility of the space, which means that if a feature  $k$  varies wildly across the dataset, a synthetic point  $x'$  may also vary this feature while remaining close to  $x_i$  under the distance metric. The use of median absolute difference rather than the more usual standard deviation also makes this metric more robust to outliers. Of equal importance are the sparsity-inducing properties of the  $L[1]$  norm. The  $L[1]$  norm is widely recognised in mathematical and machine learning circles for its tendency to induce sparse solutions in which most entries are zero when paired with an appropriate cost function.<sup>63</sup>

When computing human-understandable counterfactuals, this property is highly desirable as it corresponds to counterfactuals in which only a small number of variables are changed and most remain constant, making the counterfactuals much easier to communicate and comprehend. This metric works equally well on the examples we consider.

To demonstrate the importance of the choice of distance function, we illustrate below the impact of varying  $d(.,.)$  on the LSAT dataset. A further challenge lies in ensuring that the synthetic counterfactual  $x'$  corresponds to a valid data point. We illustrate some of the pitfalls [\*856] and remedies for dealing with discrete features when computing counterfactuals.

#### A. LSAT Dataset

We first consider the generation of counterfactuals on the LSAT<sup>64</sup> dataset. In particular, we consider a stripped-down version used in the fairness literature<sup>65</sup> that attempts to predict law students' first-year average grade on the basis of their race, grade-point average prior to law school, and law school entrance exam scores. This stripped-down version of the LSAT dataset is used in the fairness literature, as classifiers trained on this data naturally exhibit bias against 'black' people.<sup>66</sup> As a result, we will find evidence of this bias in our neural network in some of the counterfactuals we generate.

---

<sup>63</sup> See, e.g., Emmanuel J. Candes, Justin K. Romberg & Terence Tao, *Stable Signal Recovery from Incomplete and Inaccurate Measurements*, 59 COMM. PURE & APPLIED MATHEMATICS 1207, 1212 (2006).

<sup>64</sup> See R. Darrell Bock & Marcus Lieberman, *Fitting a Response Model for n Dichotomously Scored Items*, 35 PSYCHOMETRIKA 179, 187-96 (1970).

<sup>65</sup> See, e.g., Chris Russell et al., *When Worlds Collide: Integrating Different Counterfactual Assumptions in Fairness*, in ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS 6396-405 (2017); Kusner et al., *supra* note 57, at 9-12.

<sup>66</sup> See Russell et al., *supra* note 65, at 6396-405; Kusner et al., *supra* note 57, at 10.

We generate a three-layer fully connected neural network, with two hidden layers of 20 neurons each feeding into a final classifier. Even a small model like this has 941 different weights controlling its behaviour and 40 neurons that exhibit complex interdependencies, which makes conveying its internal state challenging.

Choosing  $d$  as the unweighted squared Euclidean distance,

Equation 5

we consider the counterfactual, "What would have to be changed to give a predicted score of 0?"<sup>67</sup> Directly solving for Equation 2 gives the results in the central block labelled "Counterfactuals" in Table 1. [\*857]

Unnormalised L2									
	Original Data			Counterfactuals			Counterfactual Hybrid		
score	GPA	LSAT	Race	GPA	LSAT	Race	GPA	LSAT	Race
0.17	3.1	39.0	0	3.0	39.0	0.3	1.5	38.4	0
0.54	3.7	48.0	0	3.5	47.9	0.9	-1.6	45.9	0
-0.77	3.3	28.0	1	3.5	28.1	-0.3	5.3	28.9	0
-0.83	2.4	28.5	1	2.6	28.6	-0.4	4.8	29.4	0
-0.57	2.7	18.3	0	2.9	18.4	-1.0	8.4	20.6	0

Table 1: Unnormalized  $L[2]$

Two artefacts are immediately apparent. The first is that although in this dataset, race is modelled using a discrete variable that can only take the labels 0 or 1, corresponding to 'white' or 'black' respectively, a variety of meaningless values, either fractional or negative, have been assigned to it. In the literature on adversarial perturbation, generally values are capped to lie within a sensible range such as  $[0,1]$  to stop some of these artefacts from occurring. However, this would still allow the fractional solutions shown in the top two examples. Instead, we clamp the race variable forcing it to take either value 0 or 1 in two separate run-throughs, and then take as a solution the closest counterfactual found in either of the runs. These results can be seen in the rightmost column "Counterfactual Hybrid." The algorithm now suggests always changing the race to 'white' as part of the counterfactual. Of particular note is that the counterfactuals show that 'black' students would get better scores.

The second artefact is that the algorithm much prefers significantly varying the GPA than the exam results, and this is down to our choice of distance function. We took as  $d(.,.)$ , the squared Euclidean distance, and this generally prefers changes that are as small as possible and spread uniformly across all variables. However, the range of the GPA is much smaller than that of the exam scores. Adjusting for this by normalising each component by its standard deviation, i.e.

Equation 6

gives the set of counterfactuals shown in Table 2. [\*858]

<sup>67</sup> The scores being predicted are normalised, with 0 corresponding to the average score.

Normalised L2									
score	GPA	LSAT	Race	GPA	LSAT	Race	GPA	LSAT	Race
0.17	3.1	39.0	0	3.0	37.0	0.2	3.0	34.0	0
0.54	3.7	48.0	0	3.5	39.5	0.4	3.5	33.1	0
-0.77	3.3	28.0	1	3.5	39.8	0.4	3.4	33.4	0
0.83	2.4	28.5	1	2.7	37.4	0.2	2.6	35.7	0
-0.57	2.7	18.3	0	2.8	28.1	-0.4	2.9	34.1	0

Table 2: Normalised  $L[2]$

After normalisation, the GPA remains much more consistent, and naturally remains within an expected range of values. Note that for 'black' students, race does vary under the computed counterfactual, revealing a dependence between the decision and race (which is often a legally protected attribute).

Finally, we show the use of the  $L[1]$  norm weighted by the inverse median absolute deviation (Table 3). This returns similar but sparser results to the weighted squared Euclidean distance, with the GPA not being changed under the counterfactuals.

Normalised L1									
	Original Data			Counterfactuals Continuous			Counterfactual Hybrid		
score	GPA	LSAT	Race	GPA	LSAT	Race	GPA	LSAT	Race
0.17	3.1	39.0	0	3.1	35.0	0.1	3.1	34.0	0
0.54	3.7	48.0	0	3.7	33.5	0.0	3.7	32.4	0
-0.77	3.3	28.0	1	3.3	34.4	0.1	3.3	33.5	0
-0.83	2.4	28.5	1	2.4	39.3	0.2	2.4	35.8	0
-0.57	2.7	18.3	0	2.7	35.8	0.1	2.7	34.9	0

Table 3: Normalised  $L[1]$

These final Normalised  $L[1]$  Hybrid Counterfactuals can be expressed in a more accessible text form that only describes the alterations to the original data:

**Person 1:** If your LSAT was 34.0, you would have an average predicted score (0).

**Person 2:** If your LSAT was 32.4, you would have an average predicted score (0).

**Person 3:** If your LSAT was 33.5, and you were 'white', you would have an average predicted score (0).

**Person 4:** If your LSAT was 35.8, and you were 'white', you would have an average predicted score (0).

**Person 5:** If your LSAT was 34.9, you would have an average predicted score (0).

[\*859] *B. Pima Diabetes Database*

To demonstrate Counterfactuals on a more complex problem, we consider a database used to predict whether women of Pima heritage are at risk of diabetes. <sup>68</sup> We generate a classifier that returns a risk score between  $[0, 1]$  by training a similar three layer fully connected neural-network with two hidden layers of 20 neurons to perform logistic regression. This classifier takes as input 8 different variables of varying predictive power, including number of pregnancies, age and BMI. Counterfactuals are generated to answer the question "What would have to be different for this individual to have a risk score of 0.5?" To induce sparsity in the answer and generate counterfactuals that are easy for a human to evaluate, with only a small number of changed variables, we make use of the  $L[1]$  norm, or Manhattan distance, weighted by the inverse median absolute deviation, instead of the Euclidean distance. We also cap variables to prevent them from going outside the range seen in the training data.

With this done, the counterfactuals typically vary from the original data only in a small number of variables, and these differences are automatically rendered in human readable text form.

**Person 1:** If your 2-Hour serum insulin level was 154.3, you would have a score of 0.51.

**Person 2:** If your 2-Hour serum insulin level was 169.5, you would have a score of 0.51.

**Person 3:** If your Plasma glucose concentration was 158.3 and your 2-Hour serum insulin level was 160.5, you would have a score of 0.51.

These counterfactuals are similar to the risk factors already used by doctors to communicate, e.g. "If your body mass index is greater than 40 you are morbidly obese, and at greater risk of ill-health." However, counterfactuals may make use of multiple factors and convey a personalised risk model that takes into account other attributes that may mitigate or increase risk.

### *C. Causal Assumptions and Counterfactual Explanations*

The reader familiar with causal modelling may have noticed that our counterfactual explanations are not making use of causal models [\*860] or, equivalently, that they make naive assumptions that variables are independent of one another. There are several reasons for this. One important use of counterfactual explanations is to provide the data subject with information to make a guided audit of the data and check for relevant inaccuracies in the data. Treating such errors as independent and drawn from a robust distribution such as the Laplacian (corresponding to use of the  $L[1]$  norm in our objective) is a sensible model for these errors. More importantly, creating and interpreting accurate causal models is difficult. Requiring data controllers to build and convey to a lay audience a causal model that accurately captures the interdependencies between measurements such as the number of pregnancies, age, and BMI is extremely challenging and may be irrelevant.

Counterfactuals generated from an accurate causal model may ultimately be of use to experts (e.g., to medical professionals trying to decide which intervention will move a patient out of an at-risk group). However, the purpose of our paper is to illustrate how far you can go with minimal assumptions and that such detailed causal models are unnecessary for counterfactual explanations to be of use.

## IV. ADVANTAGES OF COUNTERFACTUAL EXPLANATIONS

---

<sup>68</sup> Jack W. Smith et al., *Using the ADAP Learning Algorithm to Forecast the Onset of Diabetes Mellitus*, PROC. ANN. SYMP. ON COMPUT. APPLICATION MED. CARE 261, 261-62 (1988).

Counterfactual explanations differ markedly from existing proposals in the machine learning and legal communities (particularly regarding the GDPR's "right to explanation"),<sup>69</sup> while offering several advantages. Principally, counterfactuals bypass the substantial challenge of explaining the internal workings of complex machine learning systems.<sup>70</sup> Even if technically feasible, such explanations may be of little practical value to data subjects. In contrast, counterfactuals provide information to the data subject that is both easily digestible and practically useful for understanding the reasons for a decision, challenging them, and altering future behaviour for a better result.

The reduced regulatory burden of counterfactual explanations is also significant. Current state-of-the-art machine learning methods make decisions based upon deep networks that compose together functions more than a thousand times and with more than ten million [\*861] parameters controlling their behaviour.<sup>71</sup> As the working memory of humans can contain around seven distinct items,<sup>72</sup> it remains unclear whether "human-comprehensible meaningful information" about the logic involved in a particular decision can ever exist, let alone whether such information could be meaningfully conveyed to non-experts.<sup>73</sup> As such, regulations that require meaningful information regarding the internal logic to be conveyable to a lay audience could prohibit the use of many standard approaches. In contrast, counterfactual explanations do not attempt to convey the logic involved and, as shown in the previous Section, are simple to compute and convey.

Such expectations of providing information regarding the internal logic of algorithmic decision-making systems have surfaced recently in relation to the GDPR and in particular, the "right to explanation." The GDPR contains numerous provisions requiring information to be communicated to individuals about automated decision-making.<sup>74</sup> Significant discussion has emerged in legal and machine learning communities regarding the specific requirements and limitations of the GDPR in this regard and in particular, how to provide information about decisions made by highly complex automated systems.<sup>75</sup> As

---

<sup>69</sup> Although a right to explanation is not itself legally binding, data subjects are entitled to receive "meaningful information about the logic involved, as well as the significance and the envisaged consequences" of automated decision-making under the GDPR's Art. 13-15. Wachter, Mittelstadt & Floridi, *supra* note 1, at 16. Others have proposed that these provisions require the data subject to be given information about the internal logic and the rationale of specific decisions. The information sought aligns with the type of explanation pursued in the machine learning community. For an explanation of why such information is not legally required, and why Art. 13-15 do not constitute a de facto right to explanation, see Wachter, Mittelstadt & Floridi, *supra* note 1, at 10-11, 14-19.

<sup>70</sup> Burrell, *supra* note 5, at 9.

<sup>71</sup> See generally Kaiming He et al., *Deep Residual Learning for Image Recognition*, PROC. IEEE CONF. ON COMPUT. VISION & PATTERN RECOGNITION 770, 770-78 (2016); Gao Huang et al., *Deep Networks with Stochastic Depth*, EUROPEAN CONF. ON COMPUT. VISION 646, 646-61 (2016).

<sup>72</sup> George A. Miller, *The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information*, 63 PSYCHOL. REV. 81, 91 (1956).

<sup>73</sup> Burrell, *supra* note 5, at 9.

<sup>74</sup> Regulation 2016/679 of Apr. 27, 2016, on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation) [hereinafter GDPR], recitals 63 & 71 & arts. 13(2)(f), 14(2)(g), 15(1)(h) & 22, 2016 O.J. (L 119) 12, 14, 41, 42, 43, 46 (EU).

<sup>75</sup> See generally Wachter, Mittelstadt & Floridi, *supra* note 1; Doshi-Velez et al., *supra* note 15; Edwards & Veale, *supra* note 1; Christopher Kuner et al., *Machine Learning with Personal Data: Is Data Protection Law Smart Enough to Meet the Challenge?*, 7 INT'L DATA PRIVACY L. 1 (2017); Selbst & Powles, *supra* note 1; Malgieri & Comandé, *supra* note 1; Mendoza & Bygrave, *supra* note 1.



counterfactuals provide a method to explain some of the rationale of an automated decision while avoiding the major pitfalls of interpretability or opening the "black box," they may prove a highly useful mechanism to meet the explicit requirements and background aims of the GDPR.

## V. COUNTERFACTUAL EXPLANATIONS AND THE GDPR

Although the GDPR's "right to explanation" is not legally binding, it has nonetheless connected discussion of data protection law to the longstanding question of how algorithmic decisions can be explained to experts as well as non-expert parties affected by the decision.<sup>76</sup> [\*862] Answering this question largely depends upon the intended purpose of the explanation; the information to be provided must be tailored in terms of structure, complexity, and content with a particular aim in mind. Unfortunately, the GDPR does not explicitly define requirements for explanations of automated decision-making and provides few hints as to the intended purpose of explanations of automated decision-making.<sup>77</sup> Recital 71 of the GDPR, a non-binding provision, states that suitable safeguards against automated decisionmaking should be implemented and "should include specific information to the data subject and the right to obtain human intervention, to express his or her point of view, to obtain an explanation of the decision reached after such assessment and to challenge the decision."<sup>78</sup>

This is the only time where an explanation is mentioned in the GDPR, leaving the reader with little insight into what type of explanation is intended or what purpose it should serve. Based on the text, the only clear indication is that legislators wanted to clarify that some type of explanation can voluntarily be offered *after* a decision has been made. This can be seen as Recital 71 separates "specific information" which should be given before a decision is made,<sup>79</sup> from safeguards that apply *after* a decision has been made<sup>80</sup> ("an explanation of the decision *reached after such assessment*"<sup>81</sup>). Further indications are not provided of the intended content of such ex post explanations.<sup>82</sup>

---

<sup>76</sup> See Wachter, Mittelstadt & Floridi, *supra* note 1, at 3-4.

<sup>77</sup> See *id.* at 42.

<sup>78</sup> Regulation 2016/679, GDPR, recital 71, 2016 O.J. (L 119) 14 (EU).

<sup>79</sup> Jörg Hladjk, *DS-GVO Art. 22 Automatisierte Entscheidungen im Einzelfall einschließlich Profiling*, in DATENSCHUTZ-GRUNDVERORDNUNG 529, 535 (Eugen Ehmann & Martin Selmayr eds., 1st ed. 2017).

<sup>80</sup> The European Parliament makes the same distinction (information obligations vs. explanations of automated decisions) in their draft report on civil law rules on robotics when referring to the GDPR. See European Parliament Committee on Legal Affairs, *Draft Report with Recommendations to the Commission on Civil Law Rules on Robotics* (May 31, 2016), <http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML%2BCOMPARL%2BPE-582.443%2B01%2BDOC%2BPDF%2BV0//EN> [<https://perma.cc/A2L8-FKMP>]; Sandra Wachter, Brent Mittelstadt & Luciano Floridi, *Transparent, Explainable, and Accountable AI for Robotics*, 2 SCI. ROBOTICS 1, 1 (2017); Hladjk, *supra* note 79 at 535-36 (supporting this view that an explanation should be given after a decision has been taken, while recognising that this is not legally binding).

<sup>81</sup> Regulation 2016/679, GDPR, recital 71, 2016 O.J. (L 119) 14 (EU) (emphasis added).

<sup>82</sup> See generally European Commission, *Proposal for a Regulation of the European Parliament and the Council on the Protection of Individuals with Regard to the Processing of Personal Data and On the Free Movement of Such Data (General Data Protection Regulation)*, COM (2012) 11 final (Jan. 25, 2012), [http://ec.europa.eu/justice/data-protection/document/review2012/com\\_2012\\_11\\_en.pdf](http://ec.europa.eu/justice/data-protection/document/review2012/com_2012_11_en.pdf) [<https://perma.cc/6Q5X-F4JX>].

The content of an explanation must reflect its intended purpose. Given the lack of guidance in the GDPR, many aims for explanations [\*863] are feasible. Reflecting the GDPR's emphasis on protections and rights for individuals,<sup>83</sup> here we examine potential purposes for explanations from the perspective of the data subject. We propose three possible aims of explanations of automated decisions: to enhance understanding of the scope of automated decision-making and the reasons for a particular decision, to help contest a decision, and to alter future behaviour to potentially receive a preferred outcome. This is not an exhaustive list of potential aims of explanations, but rather reflects how the recipient of an automated decision, as with any type of decision, may wish to understand its scope, effects, and rationale and take actions in response. In the following Sections, we assess how these three purposes are reflected in the GDPR and the extent to which counterfactual explanations meet and exceed the GDPR's requirements.

#### A. Explanations to Understand Decisions

One potential purpose of explanations is to provide the data subject with understanding of the scope of automated decision-making, and the reasons that led to a particular decision. Several provisions in the GDPR can support a data subject's understanding of automated decision-making, although the types of information that must be shared tend to enhance a broad understanding of automated decisionmaking systems, as opposed to the rationale of specific decisions.<sup>84</sup> As a result, the GDPR does not appear to require opening the "black box" to explain the internal logic of the decision-making system to data subjects. With this in mind, counterfactuals can provide information aligned with the GDPR's various informational requirements, while also providing some insight into the reasons that led to a particular decision. Counterfactuals, thus, could meet and exceed the requirements of the GDPR.

The description of explanations in Recital 71 does not include a requirement to open the "black box."<sup>85</sup> Understanding the internal [\*864] logic of the algorithmic decision-making system is not explicitly required. Elsewhere, the GDPR contains transparency mechanisms,<sup>86</sup> notification duties,<sup>87</sup> and the right

---

<sup>83</sup> Christopher Kuner, *The European Commission's Proposed Data Protection Regulation: A Copernican Revolution in European Data Protection Law*, PRIVACY & SEC. L. REP. 1, 6-8 (2012).

<sup>84</sup> See Wachter, Mittelstadt & Floridi, *supra* note 1, at 5.

<sup>85</sup> *Id.*; see ARTICLE 29 DATA PROTECTION WORKING PARTY, GUIDELINES ON AUTOMATED INDIVIDUAL DECISION-MAKING AND PROFILING FOR THE PURPOSES OF REGULATION 2016/679 29 (2018), [http://ec.europa.eu/newsroom/article29/document.cfm?doc\\_id=49826](http://ec.europa.eu/newsroom/article29/document.cfm?doc_id=49826) [<https://perma.cc/8XU4-BHVG>]. The Guidelines, which are very ambiguous, seem to support the claim that such a requirement is not only absent, but also might not have been intended. On the one hand transparency in how decisions are made (Recital 71) appears to be very important. See Article 29 Data Protection Working Party, *Guidelines* at 27. However, at the same time, the guidelines state that the aim of Art. 15(1)(h) is not to create individual explanations that require understanding the internal logic of the algorithm. *Id.* Hence, the guidelines suggest that Art. 15(1)(h) calls for information about general system functionality, as is the case with its counterparts in Art. 13(2)(f) and Art. 14(2)(g). This reading of Articles 13-15 would suggest that the Article 29 Working Party does not view non-binding Recital 71 as a requirement to explain the internal logic of individual decisions, as even the legally binding text in Article 15(1)(h), which is sufficiently vague to allow such an interpretation, see Wachter, Mittelstadt & Floridi, *supra* note 1, is not thought to create such a requirement. For further support that Recital 71 does not hinge on opening the black box, see Martini, *DS-GVO Art. 22 Automatisierte Entscheidungen im Einzelfall einschließlich Profiling*, in DATENSCHUTZ-GRUNDVERORDNUNG Rn 35-37 (Boris P. Paal & Daniel A. Pauly eds., 1st ed. 2017).

<sup>86</sup> Regulation 2016/679, GDPR, art. 12, 2016 O.J. (L 119) 39-40 (EU).

<sup>87</sup> Regulation 2016/679, GDPR, arts. 13 & 14, 2016 O.J. (L 119) 40-42 (EU).

of access,<sup>88</sup> all of which create informational requirements concerning automated decision-making. Art. 13-15 describe what kind of information needs to be provided if data are collected, either immediately when collected from the data subject,<sup>89</sup> within a month when collected from a third party,<sup>90</sup> or at any time if requested from the data subject.<sup>91</sup> Among other things, Art. 12 explains how this information (as defined in Art. 13-14) should be conveyed.<sup>92</sup> Art. 12-14 suggest that data subjects must be provided with "a meaningful overview of the intended processing,"<sup>93</sup> including "the existence of automated decision-making, including profiling, referred to in Art. 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject,"<sup>94</sup> as opposed to a detailed explanation of the internal logic of a system after a decision has been made.<sup>95</sup> Rather they aim to offer a generic overview of intended processing activities, which enhances the data subject's understanding of the scope and purpose of automated decision-making.<sup>96</sup>

[\*865] Art. 12(7) clarifies that the aim of Art. 13-14 is to provide "in an easily visible, intelligible and clearly legible manner, a *meaningful overview of the intended processing*."<sup>97</sup> Two requirements are notable: (1) that the information provided must be meaningful to its recipient and broad in scope (a "*meaningful overview*"), and (2) that the notification occurs prior to processing ("*intended processing*").

To understand what would constitute a meaningful overview, the envisioned medium of disclosure is instructive. Broadly applicable information appears to be required, rather than personalised disclosures. Legal scholars have suggested that notification duties can be satisfied via updates to existing privacy

---

<sup>88</sup> Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU).

<sup>89</sup> Regulation 2016/679, GDPR, art. 13-15, 2016 O.J. (L 119) 40-43 (EU).

<sup>90</sup> Regulation 2016/679, GDPR, art. 14, 2016 O.J. (L 119) 41-42 (EU).

<sup>91</sup> Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU).

<sup>92</sup> Dirk Heckmann & Anne Paschke, *DS-GVO Art. 12 Transparente Information, Kommunikation*, in DATENSCHUTZ-GRUNDVERORDNUNG 367, 370 (Eugen Ehmann & Martin Selmayr eds., 1st ed. 2017).

<sup>93</sup> Regulation 2016/679, GDPR, art. 12(7), 2016 O.J. (L 119) 40 (EU).

<sup>94</sup> Regulation 2016/679, GDPR, arts. 13(2)(f) & 14(2)(g), 2016 O.J. (L 119) 41, 42 (EU).

<sup>95</sup> Lorenz Franck, *DS-GVO Art. 12 Transparente Information, Kommunikation*, in DATENSCHUTZ-GRUNDVERORDNUNG VO (EU) 2016/679 316, 320 (Peter Gola ed., 1st ed. 2017); Sebastian Schulz, *DS-GVO Art. 22 Automatisierte Entscheidungen im Einzelfall*, in DATENSCHUTZ-GRUNDVERORDNUNG VO (EU) 2016/679 410, 418-19 (Peter Gola ed., 1st ed. 2017); Suzanne Rodway, *Just How Fair Will Processing Notices Need to Be Under the GDPR*, 16 PRIV. & DATA PROT. 16, 16-17 (2016).

<sup>96</sup> See Kuner, *supra* note 75, at 2; ROSEMARY JAY, *GUIDE TO THE GENERAL DATA PROTECTION REGULATION: A COMPANION TO DATA PROTECTION LAW AND PRACTICE* 226 (4th Revised ed. 2017).

<sup>97</sup> Regulation 2016/679, GDPR, art. 12(7), 2016 O.J. (L 119) 40 (EU) (emphasis added).

statements or notices <sup>98</sup> (e.g., those displayed on websites or using QR codes). <sup>99</sup> This requirement does not change based on the form of data collection. <sup>100</sup> When data are collected from a third party, <sup>101</sup> an email sent to the data subject linking to the data controller's privacy statement(s) could suffice. <sup>102</sup> The same holds true for personalised links <sup>103</sup> referring to the privacy notice. One could imagine tools similar to those currently used in order to make users aware of the usage of cookies or monitoring of shopping behaviour that could satisfy the requirements in Art. 14, thus making data subjects immediately aware of data collection. <sup>104</sup> Detailed information appears to not be necessary as Art. 12(7) states that the required information can be provided along with standardised icons. <sup>105</sup> In trilogue, the European Parliament proposed several standardised icons that were ultimately not adopted (see Appendix 2). Despite this, the proposed icons reveal the initial expectations of regulators for simple, easily understood information. <sup>106</sup>

These examples suggest Art. 13-14 aim to provide a general overview of data processing that will be meaningful to all data subjects involved (e.g., all users of Twitter). The captive audience is more likely to be the general public or user base, not individual users, and [\*866] their unique circumstances. <sup>107</sup> This format of disclosure suggests notifications should be comprehensible to a general audience with mixed expertise and background knowledge. An "uneducated layperson" may be the envisioned audience for disclosures. <sup>108</sup> This coincides with the general notion of Art. 12(1) that all information and communication with the data subject has to be in a "concise, transparent, intelligible and easily accessible form," suggesting in-depth technical information and "legalese" would be inappropriate. <sup>109</sup> At a minimum, each provision

---

<sup>98</sup> See, e.g., ALAIN BENSOUSSAN, GENERAL DATA PROTECTION REGULATION: TEXTS, COMMENTARIES AND PRACTICAL GUIDELINES 113 (1st ed. 2017); JAY, *supra* note 96 at 223; Franck, *supra* note 95, at 320; Heckmann & Paschke, *supra* note 92, at 375-76; Rainer Knyrim, *DS-GVO Art. 14 Informationspflicht bei Erhebung von Daten*, in DATENSCHUTZ-GRUNDVERORDNUNG 412, 417-18 (Eugen Ehmann & Martin Selmayr eds., 1st ed. 2017).

<sup>99</sup> Lorenz Franck, *DS-GVO Art. 13 Informationspflicht bei Erhebung von Daten*, in DATENSCHUTZ-GRUNDVERORDNUNG VO (EU) 2016/679 331, 338-39 (Peter Gola ed., 1st ed. 2017).

<sup>100</sup> Article 29 Data Protection Working Party, *supra* note 85, at 25-26.

<sup>101</sup> Regulation 2016/679, GDPR, art. 14, 2016 O.J. (L 119) 41 (EU).

<sup>102</sup> Knyrim, *supra* note 98, at 415-19.

<sup>103</sup> Franck, *supra* note 95, at 322-23.

<sup>104</sup> Knyrim, *supra* note 98, at 420.

<sup>105</sup> *Id.* at 417; ARTICLE 29 DATA PROTECTION WORKING PARTY, GUIDELINES ON TRANSPARENCY UNDER REGULATION 2016/679 (2017).

<sup>106</sup> The European Commission is tasked in Art. 12(8) to develop such icons.

<sup>107</sup> See Regulation 2016/679, GDPR, recital 58, 2016 O.J. (L 119) 11 (EU); Heckmann & Paschke, *supra* note 92, at 378. Note that this information can also be provided orally. See JAY, *supra* note 96, at 216-17 (noting this also but warning that data controllers carry the burden to prove that the information was communicated).

<sup>108</sup> Franck, *supra* note 95, at 322 (noting this for the elderly, uneducated people, foreigners, or children); see also Heckmann & Paschke, *supra* note 92, at 376-77.

<sup>109</sup> JAY, *supra* note 96, at 218; Heckmann & Paschke, *supra* note 92, at 376; Article 29 Data Protection Working Party, *supra* note 85, at 25-26.

suggests that information disclosures need to be tailored to their audience, with envisioned audiences including children and uneducated laypeople.

Notifications regarding automated decision-making<sup>110</sup> face particular constraints within an overall "meaningful overview." According to the Article 29 Working Party,<sup>111</sup> the UK Information Commissioner's Office,<sup>112</sup> and other commentators,<sup>113</sup> informing the data subject about the "significance and envisaged consequences of automated decision-making" in a very simple manner, including "how profiling might affect the data subject generally, rather than information about a specific decision" will be sufficient.<sup>114</sup> For instance, an explanation of how a low rating of creditworthiness can affect payment options,<sup>115</sup> how intended data processing may result in a credit or job application being declined,<sup>116</sup> or how driving behaviour might impact insurance premiums would be sufficient.<sup>117</sup> Similarly, "meaningful [\*867] information about the logic involved" is said to require only "clarifying: of the categories of data used to create a profile; the source of the data; and why this data is considered relevant"<sup>118</sup> as opposed to a "detailed technical description about how an algorithm or machine learning works."<sup>119</sup>

This view is echoed in the Article 29 Working Party's guidelines on automated individual decision-making. First the "right to explanation" is only mentioned once in the guidelines without any further details on scope or purpose. This "right" is clearly separated from the legally binding safeguards in Art. 22(3), implying that the Article 29 Working Party sees a difference in the legal standing of Recitals and legally binding

---

<sup>110</sup> Regulation 2016/679, GDPR, art. 13(2), 14(2)(g), 2016 O.J. (L 119) 41-42 (EU).

<sup>111</sup> Article 29 Data Protection Working Party, *supra* note 85.

<sup>112</sup> Info. Comm'r's Office, *Feedback Request ? Profiling and Automated Decisionmaking* 15-16 (2017), <https://ico.org.uk/media/about-the-ico/consultations/2013894/ico-feedback-request-profiling-and-automated-decision-making.pdf> [<https://perma.cc/PC33-PU58>]. Note the UK's ICO is preparing new guidelines in the form of a living document, which will be continuously updated. See UK's Information Commissioner's Office, RIGHTS RELATED TO AUTOMATED DECISION MAKING INCLUDING PROFILING (2018), <https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/individual-rights/rights-related-to-automated-decision-making-including-profiling> [<https://perma.cc/NN7T-4E2V>].

<sup>113</sup> See, e.g., Rodway, *supra* note 95; Boris P. Paal, *DS-GVO Art. 13 Informationspflicht bei Erhebung von personenbezogenen Daten bei der betroffenen Person*, in DATENSCHUTZ-GRUNDVERORDNUNG (Boris P. Paal & Daniel A. Pauly eds., 1st ed. 2017).

<sup>114</sup> Info. Comm'r's Office, *supra* note 112, at 16.

<sup>115</sup> Paal, *supra* note 113, at Rn. 31-32; Info. Comm'r's Office, *supra* note 112, at 16.

<sup>116</sup> Rodway, *supra* note 95, at 2.

<sup>117</sup> Article 29 Data Protection Working Party, *supra* note 85, at 26.

<sup>118</sup> Info. Comm'r's Office, *supra* note 112, at 15.

<sup>119</sup> *Id.*; see also Eugen Ehmann, *DS-GVO Art. 15 Auskunftsrecht der betroffenen Person*, in DATENSCHUTZ-GRUNDVERORDNUNG, 430-31 (Eugen Ehmann & Martin Selmayr eds., 1st ed. 2017) (arguing that Art. 15 only entitles the data subject to know about the abstract logic and principles of data processing, but not the formula or code). Reference is made to Recital 63 in the English and French versions of the GDPR to support this claim. See also Article 29 Data Protection Working Party, *supra* note 85, at 25 ("The controller should find simple ways to tell the data subject about the rationale behind, or the criteria relied on in reaching the decision without necessarily always attempting a complex explanation of the algorithms used or disclosure of the full algorithm.").

provisions.<sup>120</sup> In fact, the guidelines do not even list the right to explanation in their "good practice suggestions" section.<sup>121</sup> Transparency about the fact that data controllers "are engaging in this type of activity," referring to automated decision-making, is essential and the main goal of Art. 13 and 14. The aim of these articles is thus to provide *ex ante* information.<sup>122</sup> This is also evident in the fact that the guidelines states that the phrase 'significance' and 'envisaged consequences' means "that information must be provided about intended or future processing, and how the automated decision-making might affect the data subject."<sup>123</sup> Elsewhere, the guidelines state that "details of the main characteristics considered in reaching the decision, the source of this information and the relevance" should be provided under Art 13-14.<sup>124</sup> Further, the "controller should find simple ways to tell the data subject about the rationale behind, or the criteria relied on in reaching, the decision. The GDPR requires the controller to provide meaningful information about the logic involved, not necessarily a complex explanation of the algorithms used or disclosure of the full algorithm."<sup>125</sup>

However, it must be noted that this requirement, despite referring to the decision-making rationale, seems to refer to general system [\*868] functionality rather than an explanation of an individual decision.<sup>126</sup> The guidelines state that Art 15(1)(h), which is seen to provide identical information as Art 13(2)(f) and 14(2)(g),<sup>127</sup> requires the data controller to "provide the data subject with information about the envisaged consequences of the processing, rather than an explanation of a particular decision."<sup>128</sup> This is further supported as the guidelines state that "meaningful information about the logic involved" means that "Instead of providing a complex mathematical explanation about how algorithms or machine-learning work, the controller should consider using clear and comprehensive ways to deliver the information to the data subject, for example: the categories of data that have been or will be used in the profiling or decision-making process; why these categories are considered pertinent; how any profile used in the automated decision-making process is built, including any statistics used in the analysis; why this profile is relevant to the automated decisionmaking process; and how it is used for a decision concerning the data subject."<sup>129</sup>

---

<sup>120</sup> ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85 at 27.

<sup>121</sup> *Id.* at 32.

<sup>122</sup> *Id.* at 25.

<sup>123</sup> *Id.* at 26. Further the guidelines state that data controllers can voluntary "to explain how a past decision has been made" which indicates the this is the exception to the rule.

<sup>124</sup> *Id.* at 26.

<sup>125</sup> *Id.* at 25.

<sup>126</sup> For an in-depth analysis between systems functionality and rationale of a decision see Wachter, Mittelstadt & Floridi, *supra* note 1.

<sup>127</sup> *See* ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 26.

<sup>128</sup> *Id.* at 27. "The controller should provide the data subject with general information (notably, on factors taken into account for the decision-making process, and on their respective 'weight' on an aggregate level) which is also useful for him or her to challenge the decision" is given as an example showing that only information about system functionality will be required.

<sup>129</sup> ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 31.

Overall, according to the Article 29 Working Party, the aim of Articles 13-15 is to demonstrate how automated processes help data controllers to make more accurate, unbiased, and responsible decisions and illustrate how the data, characteristics, and method used are suitable to achieve this goal.<sup>130</sup> In other words, the process of decisionmaking and the algorithm itself do not need to be fully disclosed, but rather a description of the logic of the algorithm which may include a list of data sources or variables.<sup>131</sup> This position finds further support in the Working Party's guidelines on transparency,<sup>132</sup> which [\*869] state that the notification duties in Art. 13-14 can be satisfied via standardised privacy notices, visualisation tools, and icons.

Each disclosure under Art. 13-14 must occur prior to data processing<sup>133</sup> or at the time of data collection, but before automated decision-making starts.<sup>134</sup> Evidence of this is seen in the future-oriented language used in Art. 13(2)(f) and Art. 14(2)(g),<sup>135</sup> the obligation for information about the necessity of providing data for processing,<sup>136</sup> the clarification in Art. 12(7) that information must be provided about "intended processing," the Article 29 Working Party's guidelines on transparency,<sup>137</sup> and other provisions and jurisprudence.<sup>138</sup> For automated decision-making, it is essential that information is provided ex ante, else

---

<sup>130</sup> See *id.* at 26.

<sup>131</sup> Paal, *supra* note 113, at Rn. 31-32 (seeing no difference between Art 13-15 in terms what kind of information needs to be provided); see also Paal, *DS-GVO Art. 15 Auskunftsrecht der betroffenen Person*, in *DATENSCHUTZ-GRUNDVERORDNUNG*, Rn. 31 (Paal & Pauly eds., 1st ed. 2017). Further support is offered by the text of Recital 51 proposed by the European Parliament during Trilogue, which referred to "the general logic of the data that are undergoing the processing and what might be the consequences of such processing." European Parliament Committee on Civil Liberties, Justice and Home Affairs, *Report on the Proposal for a Regulation of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data (General Data Protection Regulation)* A7-0402/2013, 21 (Nov. 21, 2013), <https://www.europarl.europa.eu/sides/getDoc.do?type=REPORT&reference=A7-2013-0402&language=EN> [<https://perma.cc/27B4-5PWC>].

<sup>132</sup> ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 115.

<sup>133</sup> Regulation 2016/679, GDPR, art. 12(7), 2016 O.J. (L 119) 40 (EU).

<sup>134</sup> See Wachter, Mittelstadt & Floridi, *supra* note 1, at 15; see also Franck, *supra* note 99; Knyrim, *supra* note 99; JAY, *supra* note 96, at 225 (arguing that the notification duties in Art. 13 need to apply *before* the data is collected); Franck, *supra* note 95, at 328 (linking this to Art. 13(2)(e) that obligates the data controllers to state "whether the provision of personal data is a statutory or contractual requirement, or a requirement necessary to enter into a contract, as well as whether the data subject is obliged to provide the personal data and of the possible consequences of failure to provide such data."). Information about the necessity to provide data must therefore be given before the data is collected. See also ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 12-13.

<sup>135</sup> See Wachter, Mittelstadt & Floridi, *supra* note 1, at 15; Frederike Kaltheuner & Elettra Bietti, *Data is Power: Towards Additional Guidance on Profiling and Automated Decision-Making in the GDPR*, 2 J. INF. RIGHTS POLICY PRACT. (2018), <https://journals.winchesteruniversitypress.org/index.php/jirpp/article/view/45> [<https://perma.cc/8YMX-EYLS>].

<sup>136</sup> Regulation 2016/679, GDPR, art. 13(2)(e), 2016 O.J. (L 119) 41 (EU).

<sup>137</sup> ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 105, at 14 states that "Articles 13 and 14 set out information which must be provided to the data subject at the commencement phase of the processing cycle."

<sup>138</sup> See Rainer Knyrim, *DS-GVO Art. 13 Informationspflicht bei Erhebung von Daten*, in *DATENSCHUTZ-GRUNDVERORDNUNG* 391, 411-12 (Eugen Ehmann & Martin Selmayr eds., 1st ed. 2017); Knyrim, *supra* note 99, at 418-19 (noting that prior notification is also in line with the Bara and others judgment of the ECJ (C-201/14; 1.10.2015) which will have major implication for the GDPR as it shows that the court views prior notification of data transfer as essential). The ruling said that when information is gathered from a third party and transferred to another data controller for further processing (e.g. based on Member State law) prior notification of the data subject -- even if no consent is

the right not to be the subject of an automated decision can never be realised. The data subject has no chance to assess the associated risks,<sup>139</sup> or whether one of Art. 22(2) grounds for allowing automated decision-making actually apply. Under Art. 22(2), automated decision-making is only lawful if it "is necessary for entering into, or performance of, a contract between the data subject and a data controller,"<sup>140</sup> if it is authorised by Member State law,<sup>141</sup> or if the data subject has given explicit consent.<sup>142</sup> If notifications do not occur [\*870] prior to processing or decision-making, data subjects would only be able to contest decisions after the fact. This can be time and cost intensive, and can result irreparable financial or reputational damage. Hence, one purpose of Art. 13-14 is to make data subjects aware of future processing<sup>143</sup> and to allow them to decide if they want their data to be processed (e.g., consent),<sup>144</sup> assess the legitimacy (based on Member State law or contract), or exercise other rights under the GDPR.<sup>145</sup>

### 1. Broader Possibilities with the Right of Access

The requirement for notification prior to processing applies only to the notification duties.<sup>146</sup> In contrast, the right of access<sup>147</sup> can be invoked at any time by the data subject, opening up the possibility of providing information available after a decision has been made (i.e., the reasons for a specific decision). However, scholars have argued that the information supplied via notification duties and the right of access is largely identical, meaning the right of access is similarly limited in terms of the scope of "meaningful information about the logic involved as well as the significance and the envisaged consequences."<sup>148</sup> Information can

---

required -- is essential. Not least because it enables the exercise of Art. 15 (right of access) and Art. 16 (right to data rectification) as soon as data is collected.

<sup>139</sup> See ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 24-25.

<sup>140</sup> Regulation 2016/679, GDPR, art. 22(2)(a), 2016 O.J. (L 119) 46 (EU).

<sup>141</sup> Regulation 2016/679, GDPR, art. 22(2)(b), 2016 O.J. (L 119) 46 (EU).

<sup>142</sup> Regulation 2016/679, GDPR, art. 22(2)(c), 2016 O.J. (L 119) 46 (EU).

<sup>143</sup> Franck, *supra* note 95, at 326-28.

<sup>144</sup> On the importance to inform the data subject accurately (e.g., risks, safeguards, rights, consequences, etc.) to allow for informed/explicit consent, see generally JAY, *supra* note 96, at 218; Heckmann & Paschke, *supra* note 92, at 374-75; Schulz, *supra* note 95, at 418-19; ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 12-13.

<sup>145</sup> See Article 29 Data Protection Working Party, *Opinion 10/2004 on More Harmonised Information Provisions* (Nov. 25, 2004), <http://www.statewatch.org/news/2004/dec/wp100.pdf> [<https://perma.cc/Y2DW-4F9Q>] (describing the need to provide meaningful information (to raise awareness) about data collection and the need to move away from long privacy statements); see also Heckmann & Paschke, *supra* note 92, at 388-89.

<sup>146</sup> Regulation 2016/679, GDPR, arts. 13 & 14 2016 O.J. (L 119) 40-42 (EU).

<sup>147</sup> Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU).

<sup>148</sup> Regulation 2016/679, GDPR, art. 15(1)(h), 2016 O.J. (L 119) 43 (EU).



thus largely be provided with identical tools (e.g., generic icons, privacy statements)<sup>149</sup> or generic templates<sup>150</sup> used for both notification and in response to access requests.

The narrower interpretation appears to be correct.<sup>151</sup> The Article 29 Working Party supports this view, explaining that the information [\*871] requirements in Art. 13(2)(f), 14(2)(g), and 15(1)(h) are identical,<sup>152</sup> while Art. 15(1)(h) requires "that the controller should provide the data subject with information about the envisaged consequences of the processing, rather than an explanation of a particular decision."<sup>153</sup> A similar argument has been made by the ICO stating that Art. 13-15 aim to "provide information about how profiling might affect the data subject generally, rather than information about a specific decision."<sup>154</sup> Additionally, the GDPR indicates a restricted scope for the right of access when compared to Art. 13-14. Personal data of other data subjects must not be disclosed, as this could infringe their privacy. Access requests can also contravene trade secrets or intellectual property rights (Art. 15(4) and Recital 63), meaning an appropriate balance between the data subject and controller's interests must be struck.<sup>155</sup>

## 2. Understanding Through Counterfactuals

Counterfactual explanations meet and exceed the aims and requirements of the GDPR's transparency mechanisms,<sup>156</sup> notification duties,<sup>157</sup> and right of access,<sup>158</sup> which provide data subjects with information to understand the scope of automated decision-making. As argued above, Recital 71 does not give any clear indication of the intended purpose or content of explanations, including whether the internal logic of the algorithm must be explained. By providing simple "if-then" statements, counterfactuals align with the requirement to communicate information to data subjects in a "concise, transparent, intelligible and easily accessible form."<sup>159</sup> They simultaneously provide greater insight into the data subject's personal

---

<sup>149</sup> Franck, *supra* note 95, at 328; Lorenz Franck, *DS-GVO Art. 15 Auskunftsrecht der Betroffenen Person*, in DATENSCHUTZ-GRUNDVERORDNUNG VO (EU) 2016/679 348, 349-50 (Peter Gola ed., 1st ed. 2017); Heckmann & Paschke, *supra* note 92, at 382-83; Paal, *supra* note 126, at Rn. 31.

<sup>150</sup> Franck argues templates will be helpful, because as soon as Art. 15 is lodged data controllers have to inform about all the information in Art. 15, regardless of the actual request. *See* Franck, *supra* note 95, at 320; *see also* Ehmann, *supra* note 119, at 431-32.

<sup>151</sup> *See* Wachter, Mittelstadt & Floridi, *supra* note 1; Michael Veale & Lilian Edwards, *Clarity, Surprises, and Further Questions in the Article 29 Working Party Draft Guidance on Automated Decision-Making and Profiling*, COMPUT. L. & SECURITY REV. (forthcoming 2018) (manuscript at 3), <https://papers.ssrn.com/abstract=3071679> (last visited May 6, 2018).

<sup>152</sup> ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 26-27.

<sup>153</sup> *Id.* at 27.

<sup>154</sup> The intention of Art. 15 is to provide a control mechanism for data subjects to request at any time more or less the same information as Art. 13-14, without having to rely on legal compliance with the notification duties by data controllers. Ehmann, *supra* note 119, at 426-27; Info. Comm'r's Office, *supra* note 112, at 16.

<sup>155</sup> Franck, *supra* note 142, at 355; Ehmann, *supra* note 119, at 434-35.

<sup>156</sup> Regulation 2016/679, GDPR, art. 12, 2016 O.J. (L 119) 39-40 (EU).

<sup>157</sup> Regulation 2016/679, GDPR, arts. 13 & 14, 2016 O.J. (L 119) 40-42 (EU).

<sup>158</sup> Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU).

<sup>159</sup> Regulation 2016/679, GDPR, art. 12(1), 2016 O.J. (L 119) 39 (EU).

situation and the reasons behind relevant automated decisions than an overview tailored to a general audience. Counterfactuals are also less likely to infringe on trade secrets or the rights and freedoms of others (e.g., privacy), since no data of other data subjects or detailed information about the algorithm needs to be disclosed, in line with restrictions on the right of access.<sup>160</sup>

**[\*872]** Perhaps most importantly, counterfactuals offer an explanation of some of the rationale of specific automated decisions, without needing to explain the internal logic of how a decision was reached (beyond a specific, limited set of dependencies between variables and the decision). This type of information is in line with the guidance mentioned above from the Article 29 Working Party<sup>161</sup> and the UK's ICO.<sup>162</sup> While opening the black box is not legally required, some information about the "logic involved" in automated decision-making must be provided.<sup>163</sup> Under the Data Protection Directive's right of access, disclosing the algorithm's source code, formula, weights, full set of variables, and information about reference groups has generally not been required.<sup>164</sup> The GDPR's right of access is likely to present similar requirements. Counterfactuals largely follow this precedent by disclosing only the influence of select external facts and variables on a specific decision. Although Art. 13(2)(f), Art. 14(2)(g), and Art. 15(1)(h) do not require information about specific decisions,<sup>165</sup> counterfactuals represent a minimal form of disclosure to inform the data subject about the "logic involved" in specific decisions. With this form of disclosure regulatory burden for data controllers is minimised, as resolving the technical difficulties of interpretability or explaining the internal logic of complex systems to non-experts is not required to compute and communicate counterfactual explanations. Counterfactuals can thus be recommended as a minimally burdensome and disruptive technique to help data subjects understand the rationale of specific decisions beyond the explicit legal requirements of Art. 13(2)(f), Art. 14(2)(g), and Art. 15(1)(h).

### *B. Explanations to Contest Decisions*

Another possible purpose of explanations is to provide information that helps contest automated decisions when an adverse or otherwise undesired decision is received. A right to contest decisions is provided as a safeguard against automated decision-making in Art. 22(3).

Contesting a decision can aim to reverse or nullify the decision and return to a status where no decision has been made, or to alter the result and receive an alternative decision. If the reasons that led to a **[\*873]** decision need explanation, the affected party can assess whether these reasons were legitimate and contest the assessment as required.

How a decision can be contested depends on whether the safeguards in Art. 22(3) (i.e., rights to obtain human intervention, express views, and contest the decision) are interpreted as a unit that must be invoked

---

<sup>160</sup> See Regulation 2016/679, GDPR, recital 63 & art. 15(4), 2016 O.J. (L 119) 12, 43 (EU).

<sup>161</sup> See generally Article 29 Data Protection Working Party, *supra* note 85.

<sup>162</sup> See generally Info. Comm'r's Office, *supra* note 112.

<sup>163</sup> See Regulation 2016/679, GDPR, arts. 13(2)(f), 14(2)(g) & 15(1)(h), 2016 O.J. (L 119) 41-43 (EU).

<sup>164</sup> For an in-depth analysis of this jurisprudence, see generally Wachter, Mittelstadt & Floridi, *supra* note 1.

<sup>165</sup> See ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 26-27.

together, or as individual rights that can be invoked separately or in any possible combination.<sup>166</sup> To gauge the scope of explanations according to their purpose and aim, different possible models for contesting an automated decision need to be assessed.

Four models are possible. If the safeguards are a unit and must be invoked together, it is likely that some human involvement is necessary to issue a new decision. This could either be a human making the decision without any algorithmic help, hence the new result is a human decision rather than an automated decision. Alternatively, a person could be required to make a decision taking the algorithmic assessment and/or the data subject's objections into account, which would be human assessment with algorithmic elements. In both cases data subjects would lose their safeguards against the subsequent decision, as both types of decision are not based "solely on automated processing" and thus do not meet the definition of automated individual decisions in Art. 22(1).<sup>167</sup> Another possibility is that a person could be required to monitor the input data and processing (e.g., based on the data subject's objections), with a new decision made solely by the algorithmic system. In this case the Art. 22(3) safeguards still apply to the new decision.<sup>168</sup> Finally, if the safeguards can be separated, and data subjects can invoke their right to contest the decision without invoking their right to obtain human intervention or express their views, a new decision could be issued with no human involvement. This decision could be contested again under Art. 22(3). It is unclear which of these models will be preferred following implementation of the GDPR.<sup>169</sup>

**[\*874]** The question remains what explanations would be helpful to contest decisions. This will depend on the contesting model. The first model where a human makes a new decision and disregards everything the algorithm suggested, an explanation of the rationale of the original decision could be informative, but will not practically impact the new decision made entirely by a human decision-maker. For each of the other models, where algorithmic involvement is envisioned, an explanation of the rationale of the decision could be helpful to identify potential grounds for contesting, such as inaccuracies in the input data, problematic inferences, or other flaws in the algorithmic reasoning.<sup>170</sup>

Even though an explanation of the rationale of a decision could be helpful to contest decisions, it does not imply an explanation is required by the GDPR or is the intended aim of the non-binding right to explanation.

---

<sup>166</sup> See Martini, *supra* note 85.

<sup>167</sup> Regulation 2016/679, GDPR, art. 22, 2016 O.J. (L 119) 46 (EU).

<sup>168</sup> However, if the new automated decision is communicated to the data subject by a person, it may not be considered "solely automated" and not subject to the Art 22(3) safeguards. The precise limitations on "solely automated" in Art 22(1) remain unclear. See also ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 10 (explaining that fabricated human involvement should not be used as a loophole).

<sup>169</sup> Either interpretation is possible, as indicated by the European Parliament's proposal to add the following text to Article 20 in an earlier draft of the GDPR, "[t]he suitable measures to safeguard the data subject's legitimate interests referred to in paragraph 2 shall include the right to obtain human assessment and an explanation of the decision reached after such assessment." Regulation 2014/0011, GDPR, art. 22, 2014 O.J. (L 119) ANNEX (EU). This text clarified that a human would need to assess the decision in question. However, this text was not adopted in the end, which leaves implementation of any of the four models possible. With that said, treating the safeguards as individually enforceable may be the most sensible option. Individuals can have an interest in expressing their views or obtaining human intervention when a decision is poorly understood or misunderstood. Both interests do not, however, necessarily lead to challenging the decision, particularly if challenges are costly or have a low likelihood of success.

<sup>170</sup> Brent Mittelstadt et al., *The Ethics of Algorithms: Mapping the Debate*, BIG DATA SOC. (2016), <http://bds.sagepub.com/lookup/doi/10.1177/2053951716679679> [<https://perma.cc/YB4Y-9MXD>].

<sup>171</sup> Recital 71 does not specify the aim of the right or what information should be revealed, and does not explicitly require the algorithm's internal logic to be explained. An explicit link is not established in the GDPR between the right to explanation and the right to contest, wherein the former would provide information necessary to exercise the latter. <sup>172</sup> Further, there is no reason to assume that the safeguards in Art. 22(3) must be exercised together, rather than independently of one another. Therefore, explanations under Recital 71 are not a necessary precondition to contest unfavourable decisions, even though this might be helpful.

Similarly, an explicit link has not been made between the right to contest and the transparency mechanisms, <sup>173</sup> notification duties, <sup>174</sup> right of access, <sup>175</sup> meaning the information provided through these rights and duties need not be explicitly tailored to help data subjects successfully contest decisions. <sup>176</sup>

**[\*875]** Nonetheless, information provided by Art. 12-15 may be helpful for contesting. Support is evident in the fact that notification duties aim to facilitate the exercise of other rights in the GDPR to increase individual control over personal data processing. <sup>177</sup> To achieve this, Art. 13(2)(b), 14(2)(c), and 15(1)(e) obligate data controllers to inform data subjects about their rights in Art. 15-21 <sup>178</sup> at the time when the data is collected, <sup>179</sup> within one month when obtained from a third party, <sup>180</sup> or at any time if requested by the data subject. <sup>181</sup> However, Art. 22 appears not to be covered by these provisions due to the odd phrasing of the obligation to inform of "the existence of automated decision-making, including profiling, referred to in Art. 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject." <sup>182</sup>

---

<sup>171</sup> See Regulation 2016/679, GDPR, recital 71, 2016 O.J. (L 119) 14 (EU).

<sup>172</sup> See Hladjk, *supra* note 79, at 535-36; Schulz, *supra* note 95, at 419-20 (arguing that "contesting" and "explaining" the decision are separate and independent safeguards).

<sup>173</sup> Regulation 2016/679, GDPR, art. 12, 2016 O.J. (L 119) 39-40 (EU).

<sup>174</sup> Regulation 2016/679, GDPR, arts. 13 & 14, 2016 O.J. (L 119) 40-42 (EU).

<sup>175</sup> Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU).

<sup>176</sup> At the same time, the Article 29 Data Protection Working Party argues that transparency in processing is essential to contesting, and that the reasons for the decisions and the legitimate basis should be known. See ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 27. However, the guidelines leave open whether this requires opening the black box and disclosing the algorithm. See *id.* This seems unlikely as the guidelines state that not even Art. 15(1)(h) aims to offer an explanation about an individual decision. See *id.* Hence it can be assumed that the information provided does not need to include an explanation of the internal logic of a specific decision.

<sup>177</sup> See Regulation 2016/679, GDPR, art. 12(2), 2016 O.J. (L 119) 40 (EU).

<sup>178</sup> See Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU). If invoked via Art. 15(1)(e), data controllers only have to inform about the rights enshrined in Art. 16, 17, 18, 19, and 21.

<sup>179</sup> Regulation 2016/679, GDPR, art. 13, 2016 O.J. (L 119) 40-41 (EU).

<sup>180</sup> Regulation 2016/679, GDPR, art. 14(3)(a), 2016 O.J. (L 119) 42 (EU).

<sup>181</sup> Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU).

<sup>182</sup> Regulation 2016/679, GDPR, art. 15(1)(h), 2016 O.J. (L 119) 43 (EU).

As argued above, Art. 13-15 will provide a meaningful overview of automated decision-making tailored to a general audience. On the surface, such an overview is not immediately useful for contesting decisions, as information about the rationale of individual decisions is not provided. In describing information to be provided about automated decision-making, these Articles explicitly refer only to Art. 22(1) and (4). It follows that data subjects do not need to be informed about the safeguards against automated decision-making such as the right to contest.<sup>183</sup> This limitation is telling. If the aim of Art. 13-15 were to facilitate contesting decisions by providing useful, individual-level information, one would expect the right to contest or Art. 22(3) to be explicitly discussed. Similarly, Art. 13-15 seem not to require informing the data subject about their right not to be subject to an automated individual decision,<sup>184</sup> from which a right to contest decisions could be inferred.<sup>185</sup> In fact, in an earlier draft of the GDPR it was suggested [\*876] that the information rights should refer to Art. 20 as a whole.<sup>186</sup> Ultimately, this approach was not adopted, suggesting that the lack of a requirement for useful information for contesting decisions was intentional.

This lack of an explicit link to the safeguards against automated decision-making is in many ways unsurprising. Art. 12-15 aim to inform data subjects about the existence of their rights in the GDPR,<sup>187</sup> and to facilitate those rights' exercise.<sup>188</sup> This does not, however, mean that the controller is required to provide other information to help the data subject to exercise her rights.<sup>189</sup> Rather, the data subject only needs to be informed about the existence of her rights, and provided with the necessary infrastructure for their exercise<sup>190</sup> (e.g., web portals for complaints), including the elimination of unnecessary bureaucratic

---

<sup>183</sup> See Regulation 2016/679, GDPR, art. 22(3), 2016 O.J. (L 119) 46 (EU); Schulz, *supra* note 95, at 419-20 (arguing that data controllers only have to inform about the safeguards after an adverse decision has been issued). In fact, Art. 12(3) introduces a very complicated model where the data controller has to inform upon request what kind of measures have been taken to satisfy a request under Art. 15-22, without being informed about the safeguards beforehand. See also, Martini, *supra* note 85, at Rn. 39-40 (also acknowledging this loophole).

<sup>184</sup> See Regulation 2016/679, GDPR, art. 22(1), 2016 O.J. (L 119) 46 (EU).

<sup>185</sup> It is important to note that Art. 13-15 obligates data controllers to inform about the general right to object to processing (which forces data controllers to stop processing) in Art. 21. Together with the information about the legitimate basis for processing provided in Art. 13(1)(c) and Art. 14(1)(c), this information could be used to contest decisions. However, this arrangement may place an unreasonable burden on the data subject. Contesting should be made as easy as possible, not least because the chances of successfully forcing the data controller to stop processing under Art. 21 are different from under Art. 22. For example, a legitimate interest of the data controller can trump a data subject's right to object to data processing under Art. 21. However, Art. 22 does not allow automated decision-making on the basis of a legitimate interest of the data controller (only explicit consent, law, or contract). This information will be useful for a data subject if they want to prevent data controllers from making decisions or to contest decisions. See Regulation 2016/679, GDPR, arts. 13-15 & 21-22, 2016 O.J. (L 119) 40-43, 45-46 (EU).

<sup>186</sup> See EUROPEAN DIGITAL RIGHTS, COMPARISON OF THE PARLIAMENT AND COUNCIL TEXT ON THE GENERAL DATA PROTECTION REGULATION 131 (2016), [https://edri.org/files/EP\\_Council\\_Comparison.pdf](https://edri.org/files/EP_Council_Comparison.pdf) [<https://perma.cc/CY6H-P97Z>].

<sup>187</sup> See Knyrim, *supra* note 98, at 415-16.

<sup>188</sup> See Heckmann & Paschke, *supra* note 92, at 371-72; Ehmann, *supra* note 119, at 425.

<sup>189</sup> See Heckmann & Paschke, *supra* note 92, at 378-79 (articulating that easily understood information provided about data subject rights in Art. 15-22 is sufficient to facilitate their exercise).

<sup>190</sup> See Bensoussan, *supra* note 98, at 114.

hurdles,<sup>191</sup> a guarantee of reasonable response time to queries lodged<sup>192</sup> as stated in Art. 12(3), and the opportunity to interact with someone who has the power to change the decision.<sup>193</sup> However, the data subject remains responsible to exercise her rights independently.<sup>194</sup> As one commentator notes, Art. 15 does not create a duty to legal consultancy;<sup>195</sup> rather, it is sufficient that the data controllers inform about the existing rights in the GDPR. Unfortunately, Recital 60, which vaguely states that "any further information necessary to ensure fair and transparent processing taking into account the specific circumstances and context in which the personal data are processed" should be provided to the data subject, does not offer additional assistance to the data subject.<sup>196</sup> This provision was [\*877] intentionally moved to the non-binding Recitals during trilogue negotiations.<sup>197</sup> Data controllers thus do not have a legal obligation to provide information that will be particularly useful for the data subject to exercise her other rights.

One final comparable restriction is notable concerning Art. 16, the right for the data subject to rectify inaccurate personal data. Data controllers are not required to specify which records most influenced a specific automated decision, which could be extremely helpful to a data subject attempting to identify inaccuracies as grounds to contest a decision. If large amounts of personal data are held, then the subject may have to check tens of thousands of items for inaccuracies.

## 1. Contesting Through Counterfactuals

Art. 13-15 thus do little to facilitate a data subject's ability to challenge automated decisions. Information is not provided about the safeguards in Art. 22(3) (e.g., the right to contest). It appears that data subjects do not need to be informed of their right not to be subject to an automated decision, which itself could imply a right to contest objectionable automated decisions. Similarly, Recital 71 has neither an explicit link to contesting decisions nor to understanding the black box. Even though an explanation could be helpful, they do not appear to be intended as a precondition for challenging decisions. If explanations were a precondition for contesting decisions, they would appear in the legally binding text. To offer greater protection to data subjects, these information gaps should be closed, meaning data controllers should inform about the right not to be subject to an automated decision and its safeguards. However, each of these seemingly intentional limitations on the information provided to data subjects suggests that information about the internal logic

---

<sup>191</sup> See Heckmann & Paschke, *supra* note 92, at 379-80.

<sup>192</sup> See Franck, *supra* note 95, at 323-24.

<sup>193</sup> Article 29 Data Protection Working Party, *supra* note 85, at 21.

<sup>194</sup> See Heckmann & Paschke, *supra* note 92, at 379-80.

<sup>195</sup> Franck, *supra* note 95, at 352.

<sup>196</sup> Regulation 2016/679, GDPR, recital 60, 2016 O.J. (L 119) 12 (EU).

<sup>197</sup> The European Commission, European Council, and European Parliament in Art. 141)(h) proposed to create a legal duty for data controllers to provide any further information beyond those in the notification duties to ensure fair and transparent data processing. See generally EUROPEAN COMMISSION, Regulation of the European Parliament and the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data (General Data Protection Regulation) (2012), [http://ec.europa.eu/justice/data-protection/document/review2012/com\\_2012\\_11\\_en.pdf](http://ec.europa.eu/justice/data-protection/document/review2012/com_2012_11_en.pdf) [<https://perma.cc/2HSG-9HX7>]; EUROPEAN DIGITAL RIGHTS, *supra* note 179, at 126-27, 129. However, this proposal was not adopted and moved to Recital 60, suggesting that there is no legal duty to provide more information than required in Art. 13-14; see Franck, *supra* note 99, at 337.

of an automated decision-making system (in compliance with "meaningful information about the logic involved as well as the significance and the envisaged consequences" in Art. 13(2)(4), 14(2)(g), and 15(1)(h)), which could facilitate contesting decisions, does not need to be provided.<sup>198</sup>

**[\*878]** Given these restrictions, counterfactuals could be helpful for contesting decisions, and thus provide greater protection for the data subject than currently envisioned by the GDPR. Regardless of the legal status of the right to explanation, the right to contest is a legally binding safeguard.<sup>199</sup> By providing information about the external factors and key variables that contributed to a specific decision, counterfactuals can provide valuable information for data subjects to exercise their right to contest. This would also be in line with the guidelines of the Article 29 Working Party, which urge that understanding decisions and knowing their legal basis is essential for contesting decisions, and is not necessarily linked to opening the black box.<sup>200</sup> An explanation that low-income led to a loan application being declined could, for example, help the data subject contest the outcome on the grounds of inaccurate or incomplete data regarding her financial situation. Understanding the internal logic of the system that led to income being considered a relevant variable in the decision, which would require a technical explanation unlike a counterfactual explanation (see Appendix 1), may be desirable in its own right, but is not absolutely necessary to contest the decision based on that variable.

Counterfactuals offer a solution and support for contesting decisions by providing data subjects with information about the reasons for a decision, without the need to open the black box. Although Art. 16 of the GDPR gives the data subject the right to correct inaccurate data used to make a decision, the data subject does not need to be informed on which data the decision depended. Where a large corpus of data has been collected, an individual without knowledge of which data is relevant or most influential on a particular decision is forced to vet all of it. This lack of information increases the burden on data subjects seeking a different outcome. Counterfactuals provide a compact and easy way to convey these dependencies (i.e., which data was influential), and to facilitate effective claims that a decision was made on the basis of inaccurate data and contest it.

### *C. Explanations to Alter Future Decisions*

From the view of the data subject, alongside understanding and contesting decisions, explanations can also be useful to indicate what could be changed to receive a desired result in the future. This purpose does not necessarily relate to the right to contest. Accurate decisions can produce unfavourable results for the data subject. The chances of successfully challenging the decision will also be low in some cases, or the costs and effort required too high. In these situations, **[\*879]** the data subject may prefer to change aspects of her situation by adapting her behaviour, and requesting a new decision once more favourable conditions exist.

Using explanations as a guide to altering behaviour to receive a desired automated decision is not directly addressed in the GDPR. This does not, however, undermine the interest data subjects have in receiving desired results from automated decision-making systems. For example, if a subject was rejected for a loan

---

<sup>198</sup> Regulation 2016/679, GDPR, art. 15(1)(h), 2016 O.J. (L 119) 43 (EU).

<sup>199</sup> Regulation 2016/679, GDPR, art. 22(3), 2016 O.J. (L 119) 46 (EU).

<sup>200</sup> See ARTICLE 29 DATA PROTECTION WORKING PARTY, *supra* note 85, at 25, 27.

due to insufficient income, a counterfactual explanation will indicate if reapplying in the event of an immediate pay raise is reasonable. The Article 29 Working Party seems to agree, stating in relevant guidelines that "tips on how to improve these habits and consequently how to lower insurance premiums" could be useful for the data subject.<sup>201</sup> For reasons outlined in Appendix 1, technical explanations that try to provide "meaningful information about the logic involved, as well as the significance and the envisaged consequences" of automated decision-making are not guaranteed to be useful in this situation.<sup>202</sup>

Counterfactuals can thus be useful for altering future decisions in favour of the data subject. By providing information about key variables and "close possible worlds" which result in a different decision, data subjects can understand which factors could be changed to receive the desired result. For decision-making models and environments with low variability over time, or models that are "artificially frozen" in time for individuals (i.e., future decisions will be made with the same model as the individual's original decision), this information can help the data subject to alter her behaviour or situation to receive her desired result in the future. Similarly, data controllers could contractually agree to provide the data subject with the preferred outcome if the terms of a given counterfactual were met within a specified period of time.

With that said, unanticipated dependencies between intentionally changed attributes and other variables, such as an increase in income resulting from a change in career, may undermine the utility of counterfactuals as guides for future behaviour. Counterfactual explanations can, however, address the impact of changes to more than one variable on a model's output at the same time. Further, regardless of the utility of counterfactuals as guidance for future behaviour, their ability to help individuals understand which data and variables were influential in specific prior decisions remains unaffected.

## [\*880] VI. CONCLUSION

We have proposed a novel lightweight form of explanation that we refer to as counterfactual explanations. Unlike existing approaches that try to provide insight into the internal logic of black box algorithms, counterfactual explanations do not attempt to clarify how decisions are made internally. Instead, they provide insight into which external facts could be different in order to arrive at a desired outcome.<sup>203</sup> Importantly, counterfactual explanations are efficiently computable for many standard classifiers, particularly neural networks. As our new form of explanation significantly differs from existing works, we have justified its nature as an explanation with reference to previous works in the philosophical literature and early A.I.

From the view of the data subject, we have assessed three purposes of explanations of automated decisions: understanding, contesting and altering. We compared these aims with the provisions of the GDPR and evaluated if they rely upon opening the black box. We concluded that the framework offers little support to achieve these goals, and does not mandate that algorithms are explainable to understand, contest or alter decisions.

---

<sup>201</sup> *Id.* at 26.

<sup>202</sup> Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU).

<sup>203</sup> The method described here is compatible with a proposal made by Citron and Pasquale to allow consumers to manually enter "hypothetical alterations" to their credit histories and view their effects. Citron & Pasquale, *supra* note 15, at 28.



The GDPR itself provides little insight into the intended purpose and content of explanations. Recital 71, the only provision that explicitly mentions explanations, does not reveal their intended purpose or content. Given the final text of the GDPR, it appears that explanations can voluntarily be offered after decisions have been made, and are not a required precondition to contest decisions. Further, there is no clear link that suggests that explanations under Recital 71 require opening the black box.

Recognising this relative lack of insight into explanations, related provisions addressing automated decision-making were examined. Notification duties defined in Art. 13-14 apply prior to data processing or before a decision is made (i.e. at the time of data collection), and provide a simple and generic overview of intended data processing activities aimed at a general audience.<sup>204</sup> This type of "meaningful overview" of automated decision-making is largely unsuitable to understand the rationale of specific decisions. Art. 13-14 similarly do not facilitate contesting decisions, owing to a lack of information to be provided about the right not to be subject to automated individual decision-making,<sup>205</sup> and its safeguards.<sup>206</sup> The right [\*881] of access<sup>207</sup> provides nearly identical information to Art. 13-14, and thus offers similarly limited value for understanding and contesting decisions. The rights and freedoms of others (e.g. privacy or trade secrets) which are protected in Art. 15(4) and Recital 63 pose an additional barrier to transparency when access requests are lodged. Across each of these Articles, technical explanations of the internal logic of automated decision-making systems are not legally mandated. Finally, offering explanations to give guidance how to receive the desired result in the future does not appear to be an aim of the GDPR, but could still be highly useful for individuals seeking alternative, more desirable outcomes.

Any future attempt to implement a legally binding right to explanation as a safeguard against automated decision-making within the GDPR framework faces several notable challenges. Automated decision-making must be based "solely on automated processing," and have "legal effects" or similarly significant effects.<sup>208</sup> Additionally, exemptions from the safeguards against automated decision-making can be introduced through Member State law.<sup>209</sup>

However, the data subject's desire to understand, contest, and alter decisions does not change based on these definitional issues. We therefore propose to move past the limitation of the GDPR and to use counterfactuals as unconditional explanations. These unconditional explanations should be given whenever requested, regardless of outcome (positive or negative decision), whether the decision was based on solely automated processes and their (legal or similar significant) effects.

Counterfactual explanations could be implemented in several ways. The transience of decision-making models suggests that counterfactuals either need to be computed automatically at the time a decision is made, or computed at a later time based on an archived copy of the model. As multiple outcomes based on

---

<sup>204</sup> Regulation 2016/679, GDPR, art. 12(7), 2016 O.J. (L 119) 40 (EU).

<sup>205</sup> Regulation 2016/679, GDPR, art. 22(1), 2016 O.J. (L 119) 46 (EU).

<sup>206</sup> Regulation 2016/679, GDPR, art. 22(3), 2016 O.J. (L 119) 46 (EU).

<sup>207</sup> Regulation 2016/679, GDPR, art. 15, 2016 O.J. (L 119) 43 (EU).

<sup>208</sup> Regulation 2016/679, GDPR, art. 22, 2016 O.J. (L 119) 46 (EU).

<sup>209</sup> Regulation 2016/679, GDPR, art. 23, 2016 O.J. (L 119) 46-47 (EU). These exemptions may be based, for example, on national security, the enforcement of civil law claims, or the protection of the data subject or the rights and freedoms of others.

changes to multiple variables may be possible, a diverse set of counterfactual explanations should be provided, corresponding to different choices of nearby possible worlds for which the counterfactual holds. These sets could be disclosed when automated decision-making occurs, or in response to specific requests lodged by individuals or a trusted third party auditor.<sup>210</sup> In any case, disclosures should occur in a reasonable window of time.<sup>211</sup>

**[\*882]** Future research should determine appropriate distance metrics and requirements for a sufficient and relevant set of counterfactuals across use sectors and cases that have very different needs. While prior philosophical debate may prove helpful, the absence of causal models in most modern classifiers, as well as the preferences of the recipient(s) of the set, must be accounted for in choosing appropriate metrics and requirements. Compared to prior discussion of measuring "closest possible worlds," setting requirements for appropriate "close possible worlds" represents a very different philosophical, social, and legal challenge.

To minimise bureaucratic burdens for data controllers and delays for data subjects and third party auditors, automated calculation and disclosure of counterfactuals would be preferable. We recommend this type of automated implementation going forward. One possible approach is to provide individuals or third party auditors with access to "auditing APIs,"<sup>212</sup> which allow users to request counterfactual explanations from the service provider, and perhaps compute them directly via the API. Access to historical decision-making models used for the decision at hand, as well as permissive terms of service that allow for such auditing, would be required.<sup>213</sup> This functionality could potentially be embedded in existing APIs.

Counterfactual explanations provide reasons why a particular decision was received (e.g., low income), offer grounds to contest it (e.g., if the data controller used inaccurate data about the income of the applicant), and provide limited "advice" on how to receive the desired results in the future (e.g., an increase of 4000 pounds/year would have resulted in a positive application). Their usage would help to resolve two primary objections to a legally binding right to explanation: first, that explaining the internal logic of automated systems to experts and non-experts alike is a highly difficult and perhaps intractable challenge; and second, that an excessive disclosure of information about the internal logic of a system could infringe on the rights **[\*883]** of others, either by revealing protected trade secrets or by violating the privacy of individuals whose data is contained in the training dataset. In contrast, counterfactuals allow an individual to receive explanations without conveying the internal logic of the algorithmic black box (beyond a limited set of dependencies), and are less likely to infringe the rights and freedoms of others than full disclosure. Assuming reasonable limitations are set on the number of counterfactuals that must be provided,

---

<sup>210</sup> See generally Wachter, Mittelstadt & Floridi, *supra* note 1.

<sup>211</sup> See, e.g., Regulation 2016/679, GDPR, art. 12(3), 2016 O.J. (L 119) 40 (EU).

<sup>212</sup> Who would shoulder the costs of hosting these APIs and computing counterfactuals is an important political issue that would require resolution. This issue goes beyond the scope of this paper. For related discussion of implementing algorithmic auditing, see Christian Sandvig et al., *Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms*, DATA & DISCRIMINATION: CONVERTING CRITICAL CONCERNS INTO PRODUCTIVE INQUIRY (May 22, 2014), <http://social.cs.uiuc.edu/papers/pdfs/JCA2014-Sandvig.pdf> [<https://perma.cc/9NKH-J69E>]; Brent Mittelstadt, *Auditing for Transparency in Content Personalization Systems*, 10 INTL. J. COMM. 12 (2016), <http://ijoc.org/index.php/ijoc/article/view/6267> [<https://perma.cc/TCN4-56QU>].

<sup>213</sup> Counterfactuals must be computed on the basis of the decision-making model at the time the decision was taken. Assuming automated decision-making models change over time, in implementations not involving automatic computation of counterfactuals at the time a decision is made (which may be cost prohibitive) it will be necessary for data controllers to keep 'audit logs' indicating the state of the decision-making model at the time of the decision.

counterfactuals are also less likely to provide information that reveals trade secrets or allows gaming of decisionmaking systems.

As a minimal form of explanation, counterfactuals are not appropriate in all scenarios. In particular, where it is important to understand system functionality, or the rationale of an automated decision, counterfactuals may be insufficient in themselves. Further, counterfactuals do not provide the statistical evidence needed to assess algorithms for fairness or racial bias. Given these limitations, more general forms of explanations and interpretability should still be pursued to increase accountability and better validate the fairness and functionality of systems.

However, counterfactuals represent an easy first step that balances transparency, explainability, and accountability with other interests such as minimizing the regulatory burden on business interest or preserving the privacy of others, while potentially increasing public acceptance of automatic decisions. Rather than waiting years for jurisprudence to dissolve all these uncertainties, we propose to abandon the narrow definitions and conditions the GDPR imposes on automated decision-making, and offer counterfactuals as unconditional explanations at the request of affected individuals.

#### [\*884] APPENDIX 1: SIMPLE LOCAL MODELS AS EXPLANATIONS

As discussed in Section II.B, 'Explanations in A.I. and Machine Learning,' approaches such as LIME<sup>214</sup> that generate simple models as local approximations of decisions make a three-way trade-off between the quality of the approximation versus the ease of understanding the function and the size of the domain for which the approximation is valid.<sup>215</sup>

To illustrate the instabilities of the approach with respect to the size of the domain, we consider a simple function of one variable. Even for problems such as this, the notion of scale, or how large a region should an explanation try to describe, is challenging with the ideal choice of scale depending on what the explanation would be used for.

As a real-world example, consider being stopped by someone in a car who asks which direction they should travel to go north. Fundamentally, this is a difficult question to answer well, with the most appropriate answer depending upon how far north they wish to travel. If they do not intend to travel far, simply pointing north gives them enough information. However, if they intend to travel further, roads that initially point north may double back on themselves or be cul-de-sacs and better directions are needed. If they intend to travel a long way, they may be better off ignoring the compass bearing entirely, and instead try to join up directly with an inter-city network.

This exact issue arises when automating explanations of decisions: the generated explanations are generic, and designed to be useful to the recipient of the explanation regardless of how they are used. However, as shown in Figure 1, the explanation -- or simplified model -- can vary wildly with the scale or range of inputs considered.

[\*885] Figure 1: Local models varying with choice of scale.

---

<sup>214</sup> See Ribeiro, Singh, & Guestrin, *supra* note 14.

<sup>215</sup> See Bastani, Kim & Bastani, *supra* note 13; Lakkaraju et al., *supra* note 47.

The grey line in each subfigure shows a local approximation of the same score curve centered at the same location in each plot. The varying range over which the approximation is computed is given by the region marked by black bars. Different choices of range e.g. top left vs. bottom left can lead to completely opposing explanations where the score either increases or decreases as the value along the bottom axis increase.

As can be seen, the direction and magnitude of the linear approximation (grey) to underlying function (black) vary dramatically with choice of domain, and deciding which approximation is most helpful to a layperson trying to understand the decision made about them is non-trivial.

To show the difficulties that would exist in either trying to use local models to either compute counterfactuals, or simply for the data subject to adjust their score, we assume that the subject desired to know how to obtain a lower score of -10 or below. In this case, none of the local approximations would be useful. The top left model, which is based on exact description of the function around point  $x$  predicts that a score of -10 would be obtained with a value of -2.5 -- corresponding to an actual score of 91.5, while the top right and lower left approximations suggest that it is not possible to obtain any score except 0, and the bottom left approximation says that -10 occurs at near 0.9 -- which actually corresponds to a local maxima.

In contrast, the counterfactual explanation for a query such as "Why was the score not below -10?" would return the answer "Because the  $x$  value was not 2.15" (the counterfactual is illustrated in [\*886] Figure 2 by the dot). Of course, it should be noted that the two approaches are generally incomparable. In much the same way, if a data subject desired to know a local linear approximation about their data point, knowledge of counterfactuals would not be helpful. However, of the two approaches, counterfactuals are the only one that will provide some indication if it is worth reapplying for a loan in the event of a pay rise.

Figure 2: Visual representation of the range of a counterfactual explanation

## [\*887] APPENDIX 2: EXAMPLE TRANSPARENCY INFOGRAPHIC

The Figure below shows several icons proposed by the European Parliament during trilogue that were ultimately not adopted as a standard. They nonetheless reveal the level of complexity expected by EU legislators when communicating information to data subjects under Art. 13-14. The reliance on generic icons suggests that individual-level, contextualised information is not required, meaning Art. 13-14 are not intended to provide a 'de facto' right to explanation comparable to the right contained in Recital 71. The relative simplicity of the icons also suggests that a broad audience is intended, comparable for example to website privacy notices. Finally, although the icons were rejected, the EC has been tasked with developing such standardised icons in the future (Art. 12(8)), meaning comparable icons are seen as an acceptable way to convey the information required by Art. 13-14.

Figure 3: Rejected icons presented during trilogue