# Deep Learning for Detecting and Explaining Unfairness in Consumer Contracts

Francesca LAGIOIA [a,b,1], Federico RUGGERI [c,1], Kasper DRAZEWSKI [a],
Marco LIPPI [d], Hans-Wolfgang MICKLITZ [a],
Paolo TORRONI [c] and Giovanni SARTOR [a,b]

[a] *European University Institute, Law Department, Italy*
[b] *CIRSFID, Faculty of Law, Alma Mater Studiorum, University of Bologna, Italy*
[c] *DISI, University of Bologna, Italy*
[d] *DISMI, University of Modena and Reggio Emilia, Italy*

**Abstract.** Consumer contracts often contain unfair clauses, in apparent violation of the relevant legislation. In this paper we present a new methodology for evaluating such clauses in online Terms of Services. We expand a set of tagged documents (terms of service), with a structured corpus where unfair clauses are liked to a knowledge base of rationales for unfairness, and experiment with machine learning methods on this expanded training set. Our experimental study is based on deep neural networks that aim to combine learning and reasoning tasks, one major example being Memory Networks. Preliminary results show that this approach may not only provide reasons and explanations to the user, but also enhance the automated detection of unfair clauses.

**Keywords.** Unfair clause detection, deep learning, memory networks

## 1. Introduction

As the pool of services existing solely in cyberspace rapidly grows, the number of online contracts concluded by clicking 'I agree' on a popup window or merely by using a given service also grows. Be it for shortage of time or information overload, most of these contracts are entered into without reading. Experiments on users reading Terms of Service (ToS) and privacy policies have indicated that users take under a minute to scroll through the contract before voicing their agreement, where it should have taken them at least 15 minutes to read [1]. Left to their own devices, consumers have neither the time nor the means to analyze every online contract they enter into, not to mention manually keeping track of any changes to the contract they are bound by. The weakness of their position becomes even more obvious when contrasted with the massive processing power and sophisticated machine learning algorithms used by businesses, tasked with collecting,

---

[1] The first two authors contributed equally to the work and are the corresponding authors. Francesca Lagioia: francesca.lagioia@eui.eu. Federico Ruggeri: federico.ruggeri6@unibo.it

processing, aggregating, and analyzing user data for the purposes of profiling, assessing risk, predicting group behavior and supporting other forms of analysis and intervention.

As the economist Ken Galbraith noted as early as in the 1950s, effective protection of consumers (and, more generally, of weaker parties) against the overbearing power of big business requires not only legal regulation and public supervision, but also the active countervailing power of consumers and their associations [2]. As the power of big business is today largely based on advanced technologies, and increasingly on AI, an effective social response also needs the support of AI [3]. In the consumer law and data protection law domains, some AI-powered applications have been developed, for instance to detect discrimination in commercial practices; extracting, categorizing and summarizing information from privacy documents; and assisting users in processing and understanding their contents [4]. A further contribution in this direction is provided by CLAUDETTE, a user-end tool and web service which uses machine learning to identify and grade potentially unfair clauses in ToS contracts. According to the Unfair Contract Terms Directive (UCTD), a "term" or "clause" is unfair if, "contrary to the requirement of good faith, it causes a significant imbalance in the parties' rights and obligations arising under the contract, to the detriment of the consumer".[2] This definition is further specified by an Annex containing an "indicative and non-exhaustive list of the terms which may be regarded as unfair" (art. 3.3) and by over 50 ECJ decisions [5].

CLAUDETTE was trained on a corpus of 50 terms of service contracts. These documents were annotated by lawyers who identified potentially unfair clauses and classified them depending on the category of unfairness. Even with this small dataset, the system was already able to achieve an average accuracy of around 80 percent in identifying potentially unfair clauses when tested on new documents [6]. The training set was later extended to 100 documents, which enabled an improvement in precision.

We are now working to enable CLAUDETTE to deal with rationales (reasons why a clause is considered unfair), for two parallel purposes: to improve its performance in detecting and classifying unfair clauses and to provide legal reasons why a clause is classified in a certain way.

The tasks of linking unfair clauses and rationales is a challenging one since the distinction between unfair and fair instances of behaviour or rules is not completely theorized. Human analysts and decision makers usually rely on their intuition, trained on their experience with relevant examples. However, humans are also able to provide explanations for their intuitions of unfairness, appealing to standards, rules and principles, possibly expressed by cases, and most significantly by judicial precedents. This capacity is usually lacking in most automated classifiers [7] available today, though a number of projects aim to improve the interpretability and the explainability of AI systems [8, 9]. Rationales are important in providing transparency and explainability, but may also play a role in learning. Contrary to the usual assumption of a conflict between performance and explanability, we will show that in some cases the acquaintance with explanations can improve the performance of a classifier. This paper follows and combines the results of our earlier work. In particular, we present a new small structured corpus, consisting of a knowledge base of rationales for the legal qualification of unfairness, used as a support for reasoning; some experimental results obtained by applying a new deep neural network model; and the extension of the classification task to a more informative classification of such clauses, now supported by forms of reasoning on context.

---

[2]See the Council Directive 93/13/EEC on Unfair Terms in Consumer Contracts, art. 3 (1).

## 2. Methodology

Legal experts can detect unfair clauses by relying on multiple sources of knowledge, such as the applicable legal regulations, the relevant judicial cases, their trained common sense. They also use these sources also for generating rationales (explanations). While a system aimed at recognizing unfair clauses cannot be expected to reason like a lawyer, it should however be expected to provide results that match the assessments that a trained lawyer would give after carefully reading the document containing such clauses.

In CLAUDETTE we have adopted a supervised machine learning approach, based on a training set of documents annotated by domain experts [6]. The system compares clauses classified as fair or unfair, and, based on such a comparison, it develops its implicit concept of unfairness. Such an approach has delivered very encouraging results, as noted above. However, the legal knowledge used by the system is restricted to the annotations (category and unfairness level) provided by the experts, which does not directly point to the rationales behind the annotations.

The aim of this work is to study whether the introduction of explicit domain knowledge, in particular a KB of rationales, can further improve the performance of our system, by enabling it to exploit rationales for unfairness in a forward-looking way. This corresponds to the idea that human lawyers use rationales not only to provide explanations for intuitions of unfairness, but also to guide such intuitions, pointing to general features relevant to unfairness, that may be shared by other similar clauses.

To provide a computable model for the forward-looking use of rationales we rely on a particular category of deep learning models, denoted as memory augmented neural networks (MANNs) [10, 11]. MANNs are a type of a recurrent neural network (RNN) that introduces an external memory block as a support for reasoning. Given an input, the model checks whether the memory contains some slot that is related to that input (e.g., through a similarity measure). Subsequently, the memory content is extracted and coupled with the given input to accomplish the classification task. Formally, we can distinguish between two phases: (i) the *memory addressing* sub-process, where the network computes some representation of the input to operate with the memory; (ii) the *reasoning* sub-process, where the new content is distilled for the resolution of the task. In our domain, the first step consists in comparing the clause to be evaluated with the relevant rationales in the memory, whereas the second step is to use the rationales to assess the category and the level of unfairness of the clause.

MANNs have been widely used for complex tasks where reasoning about the context of the given inputs plays a key role: question answering [12, 13, 14, 15], sentiment analysis [16], reading comprehension [17, 18, 19], graph analysis and navigation [20, 21]. In this paper we initially explore a simple variant [12] of the general concept of MANNs, named end-to-end memory network, since all the core operations (memory addressing and reasoning) are differentiable, and thus the model can be trained just like any other deep network. In the context of consumer contracts, we define the task of unfair clause detection as a binary classification task [6], where a given input clause can be labelled as either fair or (potentially) unfair.

The knowledge base stored in the memory consists of a fixed collection of possible rationales for unfairness. These rationales were provided by legal experts, based on their experience or on the case law (see Section 3), and linked by the same experts to the clauses they apply to. When analyzing an input clause, the system accesses the knowl-
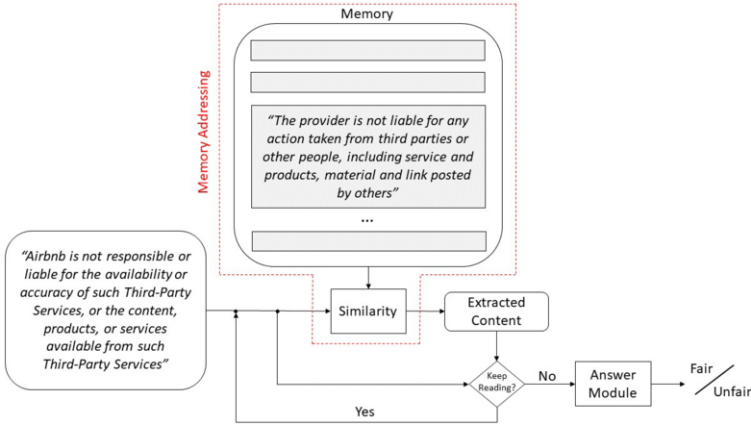
**Figure 1.** Architecture of the proposed model for unfairness detection. Each given input clause to be classified is firstly compared with the memory block via a similarity metric. By doing so, pertinent content is extracted from the memory and coupled with the input clause. Subsequently, based on the chosen fixed amount of read iterations, the model either repeats the procedure described so far with the newly modified input, or uses content gathered so far to produce a prediction via a dedicated answer module.

edge base to retrieve the rationales that best match such input. Note that the system can relate each potentially unfair statement to multiple rationales. The links between statements and rationales established by the system provide a simple criterion for qualitative analysis of the model. In particular, if clause-rationale links were given for some samples, it would be possible to compare such references made by the experts with those made by the MANN. In broad terms, the method employed by the MANN to classify a given input clause is as follows. The MANN iteratively performs reading operations based on a similarity metric between each memory slot and the input clause. Subsequently, the MANN extracts memory content by combining all memory slots, and attributing to each slot a weight proportional to the similarity score. Next, the extracted content is added to the current input to build a representation that can possibly be distilled as a new input for another iteration (reasoning phase). In this way, past iterations are always taken into account during memory reading. Eventually, after the last memory iteration, the network operates on the distilled input to predict a fair/unfair label. Note that the same memory slot may be read multiple times in order to properly exploit its content, since a distribution of weights is applied to the whole memory block. Figure 1 illustrates the architecture.

## 3. The Dataset

In our previous research[6], we produced a dataset consisting of 50 relevant online consumer contracts, i.e., Terms of Services (ToS) of online platforms. The dataset now consists of 100 ToS. Such contracts were selected among those offered by some of the major players in terms of global relevance, number of users, and time the service was established. To train the ML classifier, these ToS were analyzed and marked in XML. We focused on eight categories of clauses, which most often are unlawful or unfair, i.e., clauses establishing: (1) jurisdiction in a state different from the consumer's; (2) choice of a law

other than the consumer's; (3) limitation of provider's liability; (4) provider's right to unilaterally terminate the contract and/or access to the service; (5) provider's right to unilaterally modify the contract/service; (6) arbitration on disputes arising from the contact; (7) provider's right to unilaterally remove consumer content from the service, including in-app purchases; (8) acceptance of contract by the mere use of the service, even when the consumer has not read the contract or explicitly agreed to it [6]. As reported by [22] and as our research indicates [5], such categories are widely used in ToS for online platforms. For the purposes of this study we focused on *limitation of liability* (LTD). Clauses falling under this category stipulate that the duty to pay damages is limited or excluded for certain kinds of losses and under certain conditions. One reason for focusing on limitation of liability is that LTD is the category for which we have the largest number of problematic clauses in our dataset. More precisely, our corpus contains a total of 21,063 sentences, 674 of which contain a potentially or clearly unfair clause (note that the total number of sentences containing a potentially unfair clause, in any category, is 2,346). In particular, clauses excluding liability for broad categories of losses or causes of them were marked as potentially unfair, including those containing blanket phrases like "to the fullest extent permissible by law". Conversely, clauses meant to reduce, limit, or exclude the liability for physical injuries, intentional harm, or gross negligence were marked as clearly unfair [6, 5]. The second observation concerns the particular difficulty of detecting unfair LTD clauses. Our classifier has shown lower performance on such clauses in comparison to other categories [6]. Moreover, focusing on a single category of unfairness makes it easier to circumscribe a dedicated knowledge base for testing the MANNs. However, this does not affect the significance of our experiments, since unfair limitation of liability can be identified on the basis of several different rationales.

An initial analysis enabled us to identify 21 legal rationales for (potentially) unfair limitation of liability, which map different questionable circumstances under which the ToS reduce or exclude liability for losses or injuries. For each rationale we defined a corresponding identifier [ID]. The rationales have been formulated by two independent legal experts, each adopting different approaches. The first approach was more synthetic, and produced a smaller number of broad grounds of unfair exclusion of liability (Table 1). The second approach was more analytical, and produced many explanations, each describing multiple kinds of unfairly excluded losses or damages (Table 2). The two lists of rationales were then merged and used together for running the experiments.

Table 1.: Legal rationales for the legal qualification of unfairness (synthetic approach).

| ID | Legal Rationale |
|---|---|
| `blanket_phrase` | The limitation of liability uses a blanket phrase to the fullest extent permissible by law, any indirect or incidental damages, liability arising out of or in connection with these Terms or similar. |
| `srv_con_liab` | Liability is excluded in cases related to availability, usability or legality of service, website and/or user's content. |
| `vir_malware` | Liability is excluded for data loss, corruption or damage whether caused by viruses, trojan horses, malware or other malicious activity. |
| `physical_harm` | Liability is excluded also in cases of physical or personal injuries. |
| `third_party` | Liability is excluded for the actions and/or services of third parties. |

Table 2.: Legal rationales for the legal qualification of unfairness (analytical approach).

| ID | Legal Rationale |
|---|---|
| extent | since the clause states that to the fullest extent permissible by law the provider is not liable. |
| discontinuance | since the clause states that the provider is not liable for any technical problems, suspension, disruption, modification, discontinuance, limitation of services and features. |
| compharm | since the clause states that the provider is not liable for harm or damage to hardware and software, including viruses, worms, trojan horses, or any similar contamination or destructive program. |
| anydamage | since the clause states that the provider is not liable for any special, direct and/or indirect, punitive, incidental or consequential damage, including negligence, harm or failure. |
| amount | since the clause states that the compensation for liability or aggregate liability is limited to, or should not exceed, a certain amount. |
| thirdparty | since the clause states that the provider is not liable for any action taken from third parties or other people, including service and products, material and link posted by others. |
| security | since the clause states that the provider is not liable for any damage deriving from a security breach, including any unauthorised access. |
| disclosure | since the clause states that the provider is not liable for damages resulting from disclosure of data and personal information. |
| reputation | since the clause states that the provider is not liable for reputational and goodwill damages or loss. |
| anyloss | since the clause states that the provider is not liable for any loss resulting from the use of the service and or of the website, including lost profits, data, opportunity. |
| awareness | since the clause states that the provider is not liable whether or not he was, or should have been, aware about the possibility of any damage or loss. |
| contractfailure | since the clause states that the provider is not liable for any failure in performing contract and terms obligations, breach of agreement. |
| unilateral | since the clause states that the provider is not liable for any unilateral change or unilateral termination. |
| dataloss | since the clause states that the provider is not liable for any loss of data. |
| grossnegligence | since the clause states that the provider is not liable for gross negligence. |
| injury | since the clause states that the provider is not liable for personal injury and death. |

Each unfair limitation of liability clause in the training set has been indexed with on or more identifiers of rationales that apply to the specific clause. As an example consider the following clause taken from the terms of service of Badoo (last updated on 11 September 2018) and previously classified as potentially unfair:

```
To the fullest extent permitted by law, Badoo expressly excludes:all
conditions, representations, warranties and other terms which might otherwise
be implied by statute, common law or the law of equity; and any liability
```

```
incurred by you arising from use of Badoo, its services or these Terms,
including without limitation for any claims, charges, demands, damages,
liabilities, losses or expenses of whatever nature and howsoever direct,
indirect, incidental, special, exemplary, punitive or consequential damages
(however arising including negligence), loss of use, loss of data, loss caused
by a computer or electronic virus, loss of income or profit, loss of or damage
to property, wasted management or office time, breach of contract or claims
of third parties or other losses of any kind or character, even if Badoo has
been advised of the possibility of such damages or losses, arising out of or in
connection with the use of Badoo.
```

The clause above has been linked to the following identifiers one the analytical approach: ID: extent, anydamage, compharm, anyloss, awareness, contractfailure, dataloss. Conversely, on the synthetic approach, the clause has been associated to the following ID: blanket_phrase, vir_malware. The link between rationales and clauses will be used in future experiments to instruct the system so that it can provide an explanation for the unfairness of particular clauses.

## 4. Experimental Results

Our experiments use the proposed dataset of consumer contracts, focusing solely on LTD clauses. We employ a 10-fold cross-validation as both an evaluation and a calibration method for our models of interest. In particular, the whole corpus is first split into 10 subsets, named folds. Each fold is then used, in turn, as the test set, whereas the union of the other folds is further split into a training set and a validation set, exploited for hyper-parameter tuning. Concerning calibration, we consider two simple baselines that were also tested in some of our previous work [6] about unfairness detection in consumer contracts: (i) a network comprised of stacked recurrent neural network layers, a variant of RNNs referred to as Long Short-Term Memory network (LSTM), used extensively in the deep learning community; and (ii) a network defined as a stack of convolutional neural networks (CNNs). Moreover, we consider the current state-of-the-art solution for this task [6], featuring at its core a Support Vector Machine (SVM). Among all these models, the MANN is the only one that leverages an external knowledge base. The only input of LSTM, CNN, and SVM is the clause to be classified.

The memory network we propose follows the architecture described in [12] with minimal differences. With respect to the system illustrated in Section 2, the model performs six iterations over the memory before producing an answer. Qualitative analysis of this behaviour is reported below. Just like in [12], the similarity operation between the input clause and each content stored in memory is implemented as a simple dot product between the two sentence-embedding vectors, numerical representations of the corresponding texts. Far more complex implementations have been adopted in the literature [13, 20, 23, 24, 25], and we will consider them in future extensions. For the present study we decided to adopt the simplest similarity operation, because of the exploratory nature of our investigation. Once extracted, the distilled memory vector, defined as the weighted sum of read contents, is summed with the current input and used as input for the next iteration, as shown in Figure 1. As a last stage, a stack of fully connected layers is used to predict the classification label, given the latest memory-enhanced input.

To account for performance variations due to different initial configurations, we address model stability by repeating the cross-validation routine a sufficient amount of

**Table 3.** Results on 10-fold cross-validation. Performance measures are macro-averaged.

| Model | Precision | Recall | $F_1$ |
|---|---|---|---|
| Baseline LSTM | 37.55 | **88.93** | 51.51 |
| Baseline CNN | 43.43 | 87.90 | 56.27 |
| Memory Network | **68.36** | 84.31 | **64.33** |
| State of the art SVM | 52.52 | 81.57 | 63.7 |

times. In our experimental setting we fix the number of repetitions to 10. In this way, it is possible to gather insight about performance variance and select at the same time the best-performing results. Input sensitivity is a crucial factor for detection systems, especially when the task is centred on infrequent or hard-to-detect anomalies. Furthermore, models were early stopped based on validation loss scores. Hyper-parameters calibration was accomplished via the same evaluation method, but without repetitions.

Table 3 reports the results of the proposed experimental setting. In particular, for all the models we report precision, recall, and $F_1$ scores, macro-averaged over the ten folds.[3] From the collected results, it is evident that baseline models that exclusively leverage the input-clause content fail to correctly classify the majority of legal violations in consumer contracts. On the other hand, the MANN model shows a strong improvement in performance with respect to the baselines. This indeed corroborates the rationale that background knowledge is a crucial element for this task. The proposed model exploiting an external memory shows even slightly better results than the state-of-the-art SVM.

Differently from other described architectures, modelling explicit comparison in memory networks presents the advantage of directly visualizing the interaction level of the model with respect to its memory blocks. This opens up the possibility of understanding *what* the model believes to be useful for the detection task. As an example, Figure 2 shows the overall memory usage over all folds. It is clear that the model does not exploit all the memory slots equally. The most used memory is the one stating *"The limitation of liability uses a blanket phrase like to the fullest extent permissible by law, any indirect or incidental damages, liability arising out of or in connection with these terms, or similar"*, which is one of the most general explanations, also providing several examples. Overall, the eight most used memories account for over 45% of the cases labeled by our experts, which is very interesting, since we point out that no information about which explanations were linked to which clauses was given during training. The latter would be called a *strong* supervision for memory networks, and we plan to use it in future work.

## 5. Conclusions

This paper investigates the use of memory-augmented deep learning models, for the automated detection of potentially unfair clauses, with a focus on limitation of liability.

This study was motivated by two main goals. The short-term goal was aimed at verifying whether the use of a knowledge base of rationales can improve the system performance in unfairness detection, while improving the reliability of the classification task. Our results are very encouraging: using a relatively small set of rationales, and no information about the link between explanations and clauses, the proposed MANN markedly

---

[3]For all the neural models, we have exploited multi-start by training ten different networks for each fold, and selecting the best network for each fold according to the validation performance.
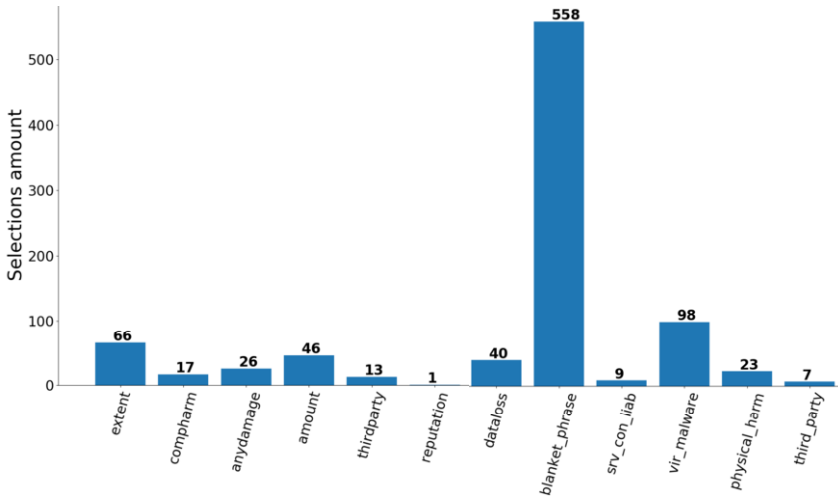
**Figure 2.** Cumulative memory distribution across all test folds. Slots never selected are not shown.

outperforms other neural models, while slightly improving over the state-of-the-art SVM. Moreover, memory-enhanced architectures inherently allow qualitative in-depth analyses of the model's behaviour, facilitating task-related investigations concerning relevant issues, such as trustworthiness and stability. Nonetheless, further steps are required, such as the annotation of memory targets for each input clause, to be exploited both to train the model towards task objectives, and to directly assess behaviour comparison. We are also working on the construction of a larger knowledge base of rationales relatively to other categories of unfair clauses, with the intention of improving training.

In the future, we plan to exploit the memory-enhanced architecture proposed in this study so as also to provide meaningful and trusted explanations for the users of CLAUDETTE, namely consumers, their organisations, and enforcement authorities.

## References

[1] J.A. Obar and A. Oeldorf-Hirsch, The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services, *Information, Communication & Society* **0**(0) (2018), 1–20.

[2] K. Galbraith, American Capitalism  The Concept of Countervailing Power, *American Journal of Agricultural Economics* **34**(4) (1952), 569–572.

[3] M. Lippi, G. Contissa, F. Lagioia, H.-W. Micklitz, P. Pałka, G. Sartor and P. Torroni, Consumer protection requires artificial intelligence, *Nature Machine Intelligence* **1**(4) (2019), 168.

[4] P. Pałka and M. Lippi, Big Data Analytics, Online Terms of Service and Privacy Policies, *Research Handbook on Big Data Law edited by Roland Vogl* (2019).

[5] H.-W. Micklitz, P. Pałka and Y. Panagis, The empire strikes back: digital control of unfair terms of online services, *Journal of consumer policy* **40**(3) (2017), 367–388.

[6] M. Lippi, P. Pałka, G. Contissa, F. Lagioia, H.-W. Micklitz, G. Sartor and P. Torroni, CLAUDETTE: An automated detector of potentially unfair clauses in online terms of service, *Artificial Intelligence and Law* **27**(2) (2019), 117–139.

[7] F. Doshi-Velez, M. Kortz, R. Budish, C. Bavitz, S. Gershman, D. O'Brien, S. Schieber, J. Waldo, D. Weinberger and A. Wood, Accountability of AI under the law: The role of explanation, *arXiv preprint arXiv:1711.01134* (2017).

[8] T. Miller, Explanation in artificial intelligence: Insights from the social sciences, *Artificial Intelligence* (2018).

[9] S. Wachter, B. Mittelstadt and C. Russell, Counterfactual Explanations without Opening the Black Box: Automated Decisions and the GPDR, *Harv. JL & Tech.* **31** (2017), 841.

[10] J. Weston, S. Chopra and A. Bordes, Memory networks, *arXiv preprint arXiv:1410.3916* (2014).

[11] A. Graves, G. Wayne and I. Danihelka, Neural Turing machines, *arXiv preprint arXiv:1410.5401* (2014).

[12] S. Sukhbaatar, J. Weston, R. Fergus et al., End-to-end memory networks, in: *Advances in neural information processing systems*, 2015, pp. 2440–2448.

[13] A. Kumar, O. Irsoy, P. Ondruska, M. Iyyer, J. Bradbury, I. Gulrajani, V. Zhong, R. Paulus and R. Socher, Ask me anything: Dynamic memory networks for natural language processing, in: *International conference on machine learning*, 2016, pp. 1378–1387.

[14] A. Bordes, Y.-L. Boureau and J. Weston, Learning end-to-end goal-oriented dialog, *arXiv preprint arXiv:1605.07683* (2016).

[15] A. Bordes, N. Usunier, S. Chopra and J. Weston, Large-scale simple question answering with memory networks, *arXiv preprint arXiv:1506.02075* (2015).

[16] D. Tang, B. Qin and T. Liu, Aspect level sentiment classification with deep memory network, *arXiv preprint arXiv:1605.08900* (2016).

[17] A. Miller, A. Fisch, J. Dodge, A.-H. Karimi, A. Bordes and J. Weston, Key-value memory networks for directly reading documents, *arXiv preprint arXiv:1606.03126* (2016).

[18] F. Hill, A. Bordes, S. Chopra and J. Weston, The Goldilocks principle: Reading children's books with explicit memory representations, *arXiv preprint arXiv:1511.02301* (2015).

[19] J. Cheng, L. Dong and M. Lapata, Long short-term memory-networks for machine reading, *arXiv preprint arXiv:1601.06733* (2016).

[20] A. Graves, G. Wayne, M. Reynolds, T. Harley, I. Danihelka, A. Grabska-Barwińska, S.G. Colmenarejo, E. Grefenstette, T. Ramalho et al., Hybrid computing using a neural network with dynamic external memory, *Nature* **538**(7626) (2016), 471.

[21] W. Zaremba and I. Sutskever, Reinforcement learning neural turing machines-revised, *arXiv preprint arXiv:1505.00521* (2015).

[22] M. Loos and J. Luzak, Wanted: a bigger stick. On unfair terms in consumer contracts with online service providers, *Journal of consumer policy* **39**(1) (2016), 63–90.

[23] C. Xiong, S. Merity and R. Socher, Dynamic memory networks for visual and textual question answering, in: *International conference on machine learning*, 2016, pp. 2397–2406.

[24] M. Henaff, J. Weston, A. Szlam, A. Bordes and Y. LeCun, Tracking the world state with recurrent entity networks, *arXiv preprint arXiv:1612.03969* (2016).

[25] J. Pavez, H. Allende and H. Allende-Cid, Working memory networks: Augmenting memory networks with a relational reasoning module, *arXiv preprint arXiv:1805.09354* (2018).