Fake news through influence models: a proposal

Davide Fazio[1], Pierluigi Graziani[2], Mirko Tagliaferri[3], and Raffaele Mascella[1]

[1] University of Teramo, Italy
[2] University of Urbino Carlo Bo, Italy
[3] TU Wien, Austria

Abstract. The rapid spread of digital media has facilitated the emergence and circulation of fake news, raising urgent concerns about their social, political, and economic impact. In this research paper we propose a formal logical framework to investigate the concept of a news and of a fake news based on the concept of influence within a network of agents characterized by their beliefs and shares at different time steps. Specifically, our approach builds on the notion of conditional influence, which captures the causal dependencies between agents' declarations and beliefs, and provides the basis for a precise definition of news. A news is characterized not merely as a statement, but as a statement that produces observable effects on others, as witnessed by the propagation of beliefs or further declarations. By formalizing these mechanisms, we aim to complement empirical and computational studies on fake news with an abstract and rigorous tool for reasoning about influence and diffusion.

1 Introduction

The rapid diffusion of digital media has radically transformed the ways in which information is produced, disseminated, and consumed. While these dynamics have enabled unprecedented levels of connectivity and participation, they have also facilitated the rise of *fake news*, namely, intentionally fabricated or misleading information presented as fact. Fake news poses a serious threat to the integrity of public discourse, undermines trust in legitimate sources, and can exert tangible influence on social, political, and economic decision-making processes. As recent events have demonstrated, fake news can propagate faster than corrective information, thereby amplifying its societal impact[4].

Despite the growing attention from both policymakers and the scientific community, the phenomenon of fake news remains difficult to analyse rigorously.

[4] It is worth clarifying the terminology adopted in this work. In the literature, a distinction is commonly made between *misinformation* and *disinformation*. Misinformation refers to false or misleading information that is spread without the intent to deceive, for instance, when users unknowingly share incorrect content. Disinformation, instead, designates intentionally fabricated or manipulated content aimed at deceiving or manipulating the audience, often as part of strategic campaigns. Throughout this paper, we use the term *fake news* as an umbrella notion that encompasses both misinformation and disinformation, while retaining the possibility to distinguish between them when the context requires.

Most existing approaches focus on descriptive or empirical analyses, often relying on statistical correlations [13], network structures [6], or linguistic patterns [9]. Alongside these works, there are other studies concerned with the development of logical frameworks for the analysis of the mechanisms that govern the emergence and spread of fake news. [1,7] proposes a logic-based framework using News-Diffusion Temporal Logic (NDTL) to analyse how fake and true news circulate over time. The method generates logical formulas that capture temporal propagation patterns, offering transparency and interpretability in classification. As the author shows, experimental evaluations on news diffusion datasets show the framework can effectively differentiate fake from true news based on their diffusion dynamics.
The contribution of this work aims to enrich the literature on logic-based approaches to the study of news and fake news by introducing a framework grounded in the concept of influence within a social network of agents, each endowed with specific beliefs and engaged in sharing information at given temporal steps. In particular, we represent the concept of news and fake news in a social network by means of a simple modification of the causal models introduced by Halpern [10]. In particular, a proposition φ is considered news if it is shared by a member a within a group of agents at a given time step t and exerts influence on at least one other member at a subsequent time. Consistently, a fake news is represented as a news φ disseminated by an agent a at a given time step t, which is false with respect to the state of the world at t (and at any $t' \geq t$) and, moreover, either the agent is aware of the falsity of φ, or shares φ without any concern for its truthfulness or falsity. In our framework, this latter case is formalised by attributing to a, at the time step t when φ is formulated, a belief that contains neither φ nor its negation.

Each agent is modelled by two endogenous variables: one representing its belief set, and another representing the information it shares.
As in causal models, the interactions between agents' beliefs and shares are governed by (in our framework, time-dependent) structural equations, whose purpose is to describe how, at each temporal step t, an agent's beliefs and shares depend on the current "state of the world" at that step, on its current beliefs, and on the communications of other agents. The appeal of a structural-equation-based approach lies not only in its permeability to concepts and methods already applied in sociology (see, e.g., [3,16]), but also in its ability to investigate how the very notions of news and fake news may evolve as assumptions vary (e.g., exposure, specific cognitive heuristics for belief revision, or dependencies on the shares of particular agents) upon which the equations can be designed. As a final output, we aim at providing a formalisation of human reasoning regarding how news and fake news can be generated, disseminated, and countered within a multi-agent environment. Such a formal approach is not intended to replace empirical or computational models, but rather to complement them: it offers an abstract and logically precise layer that can be integrated with other methodologies for analysing the dynamics of fake news.

The article is structured as follows. Section 2 presents a formal logical framework to investigate the concept of news and of fake news based on the concept of influence within a network of agents characterised by their beliefs and shares, inspired by ideas given in [2]. Section 3 summarises the contributions of this work and outlines possible directions for future research.

2 Formal Framework

Studies on fake news use many different methods, reflecting the fact that the problem touches both social sciences and computer science [1]. Reviews of the literature [4,12,17] show that research in social sciences often use *surveys* and *experiments* to measure people's exposure to fake news and its influence on opinions or behaviour, while computer scientists rely more on *data analysis* and *machine learning* to detect misleading content automatically. Buitrago López et al. [12], for instance, identify three main lines of research: *social frameworks* focusing on actors and contexts, approaches that model the spread of fake news as if it were an *epidemic*, and approaches that simulate how people update their beliefs when exposed to false information.

Our work is intended to complement these perspectives by introducing a formal logical framework that enables rigorous reasoning about fake news. Rather than providing empirical measurements or simulations, our aim is to offer an abstract system in which the causal mechanisms underlying the creation, dissemination, and persistence of fake news can be represented and analyzed. In this sense, the framework provides a formal foundation that can be integrated with empirical and computational approaches to enrich the overall understanding of the phenomenon.

Before introducing our formal framework, it is worth recalling that there exists a substantial line of research on causality in social networks based on structural equation models (SEMs), most notably developed by Halpern [2]. These approaches provide a rigorous account of actual causality and have been fruitfully applied to capture dependencies among variables representing social and informational phenomena. Our proposal draws inspiration from this tradition: it is not intended to replace SEMs, but to complement them with a logical perspective. While SEMs offer a statistical and computational toolbox for causal inference, our framework provides a logic formal setting in which these concepts can be rigorously reasoned about and systematically analyzed.

Let Ag be a non-empty finite set of agents. From now on, we assume as the "underlying" logic of our framework Classical Logic $\mathbf{CL} = \langle \mathbf{Fm}_{\mathcal{L}}, \vdash_{\mathbf{CL}} \rangle$, in the customary language $\{\neg, \vee, \wedge, \rightarrow\} = \mathcal{L}$. $\mathbf{Fm}_{\mathcal{L}}$ will denote the absolutely free algebra of formulas over $\mathcal{L}$ generated by an infinite countable set $Var = \{p_0, \dots, p_n, p_{n+1}, \dots\}$ of variables with universe $\mathrm{Fm}_{\mathcal{L}}$. Moreover, for any set A, $\wp(A)$ will stand for the usual powerset of A.

Definition 1 (Signature). *A* signature *for* Ag *is a tuple* $\mathfrak{S} = (\Delta, \Upsilon, r)$*, where:*

1. $\Delta = \{D_i\}_{i \in \mathsf{Ag}}$ *is a finite set of endogenous variables for* agents' declarations*;*
2. $\Upsilon = \{B_i\}_{i \in \mathsf{Ag}}$ *is a finite set of endogenous variables for* agents' beliefs*;*
3. $r : \Delta \cup \Upsilon \to \wp(\wp(\mathrm{Fm_{CL}}))$ *such that, for any* $i \in \mathsf{Ag}$, $|x| = 1$, *for any* $x \in r(D_i)$.

A signature describes the endogenous variables involved in a model, which can be affected by interactions between agents. In the next definition, $\mathrm{Th}_{\vdash_{\mathbf{CL}}}$ stands for the set of subsets of $\mathrm{Fm}_{\mathcal{L}}$ closed under $\vdash_{\mathbf{CL}}$, namely theories of Classical logic, and $\mathrm{Th}^{\max}_{\vdash_{\mathbf{CL}}}$ will stand for the set of its maximal elements which, as it is well known, are into one to one correspondence with Boolean evaluations.

Definition 2 (History). *A* history *is a function* $h : \mathbb{N} \to \mathrm{Th}^{\max}_{\vdash_{\mathbf{CL}}}$.

A history is nothing but a sequence of sets of formulas (which are maximal **CL**-theories) formalising the set of propositions describing the state of the world at each time step. As the reader may observe, the assumption that the set of true propositions at each time step constitutes a maximal theory of classical logic may appear to be a rather strong theoretical postulate. However, what we intend to represent here is not what is known by the agents (or, more generally, by human beings). Rather, it is the description of the *entire* state of affairs (known and unknown) – the *state of the world* – in which communication and belief revision take place. It is designed to ensure that the structural equations governing agents' beliefs and shares may depend on conditions external to, and independent of, the agents themselves.

Influence models are aimed at formalising a system evolving from an initial time t_0 to $+\infty$ involving agents, their beliefs and declarations (sharings), and outlining how they interact and evolve at any step of time. Our models will be completely deterministic. Indeed, given the specification of initial beliefs and sharings of agents, and of structural equations describing how each agent updates her beliefs at time t according to previous beliefs, the current state of the world, and sharings of other agents, and shares information (at time t) according to the state of the world and belief at t, and previous sharings from other agents. The deterministic nature of our model ensures a rigorous formalisation of news and fake news propagation, given fixed initial states and structural equations. We acknowledge that this oversimplifies the inherently stochastic dynamics of information diffusion, where randomness, noise, and individual variability play a role. Future work could extend the framework to include probabilistic transmission, noisy perception, or agent-based simulations to capture such uncertainties. Each model serves to describe a possible evolution of the system in a deterministic manner. Nonetheless, nothing prevents multiple models from being considered simultaneously, thereby restoring a more complex and accurate representation of the potential evolutions of the system's state.

Let $\mathfrak{S} = (\Delta, \Upsilon, r)$ be a signature. A family of *structural equations* for beliefs is a family $\{f_i^B\}_{i \in \mathsf{Ag}}$ of functions such that, for any $i \in \mathsf{Ag}$,

$$f_i^B : \mathbb{N} \times \wp(\mathrm{Fm}_{\mathcal{L}}) \times r(B_i) \times (\Pi_{j \in \mathsf{Ag} \smallsetminus \{i\}} r(D_i)) \to r(B_i),$$

and $f_i^B(0,-)$ is *constant*. In what follows we denote $f_i^B(0,-)$ by In_i^B. Similarly, we define a family of structural equations for declarations as a family $\{f_i^D\}_{i\in\mathsf{Ag}}$ of functions such that, for any $i \in \mathsf{Ag}$:

$$f_i^D : \mathbb{N} \times \wp(\mathrm{Fm}_{\mathcal{L}}) \times r(B_i) \times (\Pi_{j\in\mathsf{Ag}\smallsetminus\{i\}} r(D_i)) \to r(D_i),$$

and $f_i^D(0,-)$, henceforth denoted by In_i^D, is constant.

It is important to note that, as we have defined them, the structural equations may depend not only on the values of the B_i's, of the D_i's, and on the set of formulas associated with the current knowledge of the world, but also on time. This allows us to assign to the agents a time-dependent behaviour with respect to their beliefs and their statements, and therefore not completely determined by their beliefs and statements alone. In this way, we can take into account the fact that, at different times, each agent may be differently exposed to the statements of other agents.

Definition 3 (Influence model). *Given a signature $\mathfrak{S} = (\Delta, \Upsilon, r)$ for Ag, a history h and families of structural equations for beliefs and declarations $\{f_i^B\}_{i\in\mathsf{Ag}}$ and $\{f_i^D\}_{i\in\mathsf{Ag}}$, respectively, we call the tuple*

$$\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i\in\mathsf{Ag}}, \{f_i^D\}_{i\in\mathsf{Ag}})$$

an influence model *for* Ag.

Given an influence model $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i\in\mathsf{Ag}}, \{f_i^D\}_{i\in\mathsf{Ag}})$ for Ag, we define, for each $i \in \mathsf{Ag}$, its *belief* and *declaration* histories induced by $\mathcal{M}$ as follows.

Intuitively, a belief (declaration) history for an agent i induced by an influence model $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i\in\mathsf{Ag}}, \{f_i^D\}_{i\in\mathsf{Ag}})$ is nothing but a function $H_{B_i}^{\mathcal{M}} : \mathbb{N} \to r(B_i)$ ($H_{D_i}^{\mathcal{M}} : \mathbb{N} \to r(D_i)$) describing the evolution of beliefs of agent i depending on the initial state In_i^B (In_i^D) and defined, for each time step t, according to the structural equation f_i^B (f_i^D) that links beliefs (declarations) of i at t with her beliefs and declarations of other agents occurred in the preceding time step. Indeed, we define $H_{B_i}^{\mathcal{M}}$ and $H_{D_i}^{\mathcal{M}}$ by *mutual induction* as follows:

$$\begin{aligned} H_{B_i}^{\mathcal{M}}(0) =& \mathrm{In}_i^B \\ H_{B_i}^{\mathcal{M}}(t+1) =& f_i^B(t+1, h(t+1), H_{B_i}^{\mathcal{M}}(t), \Pi_{j\in\mathsf{Ag}\smallsetminus\{i\}} H_{D_j}^{\mathcal{M}}(t)) \end{aligned}$$

$$\begin{aligned} H_{D_i}^{\mathcal{M}}(0) =& \mathrm{In}_i^D \\ H_{D_i}^{\mathcal{M}}(t+1) =& f_i^D(t+1, h(t+1), H_{B_i}^{\mathcal{M}}(t+1), \Pi_{j\in\mathsf{Ag}\smallsetminus\{i\}} H_{D_j}^{\mathcal{M}}(t)) \end{aligned}$$

At each time step, agents revise their beliefs according to their beliefs and declarations of other agents at the previous time step, and provide a declaration according to their current beliefs and shares of other agents at the previous time step. Note that structural equations for declarations include not only agents' beliefs but also past declarations of other agents. This provides an account of the fact that, in some cases, declarations do not depend solely on agents' beliefs but also on what other agents share. For example, an agent A might be keen to declare, according to her beliefs, that the Earth is flat. However, since no other agent shares the same information, A might refrain from asserting it.

Our framework is not solely aimed at describing how agents' beliefs and statements evolve over time, but also at reasoning about how, at a given time t, their statements and beliefs would have been if, at a previous time step, things had "gone differently".

Let $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i\in\mathsf{Ag}}, \{f_i^D\}_{i\in\mathsf{Ag}})$ be an influence model for Ag, $\Gamma \in\ r(B_i)$ $\{\varphi\} \in r(D_i)$, $i \in \mathsf{Ag}$. For $t \geq 0$ let us define the model

$$\mathcal{M}^{B_i\leftarrow\Gamma|t} = (\mathfrak{S}, h, \{f_j^B\}_{j\in\mathsf{Ag}\smallsetminus\{i\}} \cup \{{}^{B_i\leftarrow\Gamma|t}f_i^B\}, \{f_j^D\}_{j\in\mathsf{Ag}}, f_0^B, f_0^D),$$

where

$$ {}^{B_i\leftarrow\Gamma|t}f_i^B(n,-) = \begin{cases} f_i^B(n,-) & \text{if } n \neq t \\ \Gamma, & \text{otherwise.} \end{cases} \tag{1}$$

Similarly, we define the model $\mathcal{M}^{D_i\leftarrow\varphi|t}$ as the influence model

$$\mathcal{M}^{D_i\leftarrow\varphi|t} = (\mathfrak{S}, h, \{f_j^B\}_{j\in\mathsf{Ag}}, \{f_j^D\}_{j\in\mathsf{Ag}\smallsetminus\{i\}} \cup \{{}^{D_i\leftarrow\varphi|t}f_i^B\}),$$

where

$$ {}^{D_i\leftarrow\varphi|t}f_i^D(n,-) = \begin{cases} f_i^D(n,-) & \text{if } n \neq t \\ \{\varphi\}, & \text{otherwise.} \end{cases} \tag{2}$$

A model $\mathcal{M}$ having e.g. the form $\mathcal{M}'^{B_i\leftarrow\Gamma}$ will be called a *counteractual model* of $\mathcal{M}'$. The following proposition is easy to prove. Therefore, it is left to the reader.

Proposition 1. *Let* $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i\in\mathsf{Ag}}, \{f_i^D\}_{i\in\mathsf{Ag}})$ *be an influence model for* Ag, $i, j \in \mathsf{Ag}$, $\Gamma_1, \Gamma_2 \in\ r(B_i)$, $\Theta \in R(B_j)$, $\{\varphi\}, \{\psi\} \in r(D_i)$, $\{\chi\} \in r(D_j)$, $t, t' \in \mathbb{N}$. *The following hold:*

1. *If* $t = t'$ $(\mathcal{M}^{B_i\leftarrow\Gamma_1|t'})^{B_i\leftarrow\Gamma_2|t} = \mathcal{M}^{B_i\leftarrow\Gamma_2|t}$;
2. *If* $t = t'$, $(\mathcal{M}^{D_i\leftarrow\varphi|t'})^{D_i\leftarrow\psi|t} = \mathcal{M}^{D_i\leftarrow\psi|t}$;
3. *If* $t \neq t'$, $(\mathcal{M}^{B_i\leftarrow\Gamma_1|t})^{B_i\leftarrow\Gamma_2|t'} = (\mathcal{M}^{B_i\leftarrow\Gamma_2|t'})^{B_i\leftarrow\Gamma_1|t}$;
4. *If* $t \neq t'$, $(\mathcal{M}^{D_i\leftarrow\varphi|t'})^{D_i\leftarrow\psi|t} = (\mathcal{M}^{D_i\leftarrow\psi|t})^{D_i\leftarrow\varphi|t'}$;
5. $(\mathcal{M}^{B_i\leftarrow\Gamma_1|t'})^{D_j\leftarrow\chi|t} = (\mathcal{M}^{D_j\leftarrow\chi|t})^{B_i\leftarrow\Gamma_1|t'}$;
6. *If* $i \neq j$, $(\mathcal{M}^{B_i\leftarrow\Gamma_1|t'})^{B_j\leftarrow\Theta|t} = (\mathcal{M}^{B_j\leftarrow\Theta|t})^{B_i\leftarrow\Gamma_1|t'}$;
7. *If* $i \neq j$, $(\mathcal{M}^{D_i\leftarrow\varphi|t'})^{D_j\leftarrow\chi|t} = (\mathcal{M}^{D_j\leftarrow\chi|t})^{D_i\leftarrow\varphi|t'}$.

In view of Proposition 1, whenever $X_1, \ldots, X_n$ are distinct variables, $x_i \in r(X_i)$, $t_1, \ldots, t_n \in \mathbb{N}$, we will write $\mathcal{M}^{[\overrightarrow{X}\leftarrow\overrightarrow{x}|\overrightarrow{t}]}$ instead of

$$(\cdots((\mathcal{M}^{X_1\leftarrow x_1|t_1})^{X_2\leftarrow x_2|t_2})^{\cdots})^{X_n\leftarrow x_x|t_n}.$$

The following example, while not especially significant from a theoretical standpoint, serves to illustrate an influence model.

Example 1. Let $\mathsf{Ag} = \{a, b\}$. Let $\mathcal{M}$ be the model defined upon setting $\mathfrak{S} = (\{B_a, B_b\}, \{D_a, D_b\}, r)$, where $r(B_x) = \wp(\mathrm{Fm}_{\mathcal{L}})$ and $r(D_x) = \{\{\varphi\} : \varphi \in \mathrm{Fm}_{\mathcal{L}}\}$, for any $x \in \mathsf{Ag}$. Moreover, let $h : \mathbb{N} \to \wp(\mathrm{Fm}_{\mathcal{L}})$ be the constant function such that $h(t) = Cn_{\vdash_{\mathbf{CL}}}(\{p, q\})$ (for all $t \in \mathbb{N}$), for some fixed propositional variables p, q. Furthermore, for any $n > 0$, we set

$$f_x^B(n, X, Y, \varphi) = \begin{cases} Cn_{\vdash_{\mathbf{CL}}}(Y, \varphi) & \text{if } Cn_{\vdash_{\mathbf{CL}}}(X, \varphi) \neq \mathrm{Fm}_{\mathcal{L}},\ Cn_{\vdash_{\mathbf{CL}}}(Y, \varphi) \neq \mathrm{Fm}_{\mathcal{L}} \\ Y & \text{otherwise} \end{cases}$$

and $f_a^B(0, -) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\})$ and $f_b^B(0, -) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q\})$. Finally, for any $x \in \mathsf{Ag}$, we set, for any $n > 0$,

$$f_x^D(n, X, Y, \varphi) = \begin{cases} \{\neg\varphi\} & \text{if } Cn_{\vdash_{\mathbf{CL}}}(X, \varphi) = \mathrm{Fm}_{\mathcal{L}} \text{ or } Cn_{\vdash_{\mathbf{CL}}}(Y, \varphi) = \mathrm{Fm}_{\mathcal{L}} \\ \top & \text{otherwise} \end{cases}$$

where $\top$ is an arbitrary tautology, and $f_a^D(0, -) := \{p\}$ and $f_b^D(0, -) := \{q\}$. This model induces the following belief and declaration histories:

$$\begin{array}{ll} H_{B_a}(0) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\}) & H_{D_a}(0) = \{p\} \\ H_{B_a}(1) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\}) & H_{D_a}(1) = \{\neg q\} \\ H_{B_a}(2) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\}) & H_{D_a}(2) = \{\neg\neg p\} \\ H_{B_a}(3) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\}) & H_{D_a}(3) = \{\neg\neg\neg q\} \\ \vdots & \vdots \end{array}$$

$$\begin{array}{ll} H_{B_b}(0) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q\}) & H_{D_b}(0) = \{q\} \\ H_{B_b}(1) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q\}) & H_{D_b}(1) = \{\neg p\} \\ H_{B_b}(2) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q\}) & H_{D_b}(2) = \{\neg\neg q\} \\ H_{B_b}(3) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q\}) & H_{D_b}(3) = \{\neg\neg\neg q\} \\ \vdots & \vdots \end{array}$$

If one considers the model $\mathcal{M}^{D_a \leftarrow (q \to r)|1}$, one obtains the following belief and declaration histories.

$$\begin{array}{ll} H_{B_a}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(0) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\}) & H_{D_a}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(0) = \{p\} \\ H_{B_a}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(1) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\}) & H_{D_a}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(1) = \{q \to r\} \\ H_{B_a}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(2) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\}) & H_{D_a}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(2) = \{\neg\neg p\} \\ H_{B_a}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(3) = Cn_{\vdash_{\mathbf{CL}}}(\{p, \neg q, \neg r\}) & H_{D_a}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(3) = \{\top\} \\ \vdots & \vdots \end{array}$$

$$
\begin{array}{ll}
H_{B_b}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(0) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q\}) & H_{D_b}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(0) = \{q\} \\
H_{B_b}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(1) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q\}) & H_{D_b}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(1) = \{\neg p\} \\
H_{B_b}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(2) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q, r\}) & H_{D_b}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(2) = \{\top\} \\
H_{B_b}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(3) = Cn_{\vdash_{\mathbf{CL}}}(\{\neg p, q, r\}) & H_{D_b}^{\mathcal{M}^{D_a \leftarrow (q \to r)|1}}(3) = \{\neg\neg\neg p\} \\
\vdots & \vdots
\end{array}
$$

As illustrated in Example 1, our formal framework allows us to represent a possible evolution of the system of a set of potentially interconnected agents, by specifying their initial beliefs and declarations and recursively deriving their subsequent belief states and declarations. What we find particularly interesting in the present framework is its ability to specify the criteria according to which agents revise their beliefs, issue declarations (possibly) in accordance with them and on the basis of further criteria, and interact with one another (depending on how the structural equations are defined). Of course, each model should be regarded as one among many possible descriptions of the system's evolution. A more accurate account ought to consider a range of possible models of the system's evolution, so as to take into account non-deterministic aspects that a single model alone cannot capture.

In the sequel we will be considering a fixed signature $\mathfrak{S} = (\Delta, \Upsilon, r)$, where $\Delta := \{B_i\}_{1 \leq y \leq n}$ and $\Upsilon := \{D_i\}_{1 \leq y \leq n}$, for some $n > 1$. Let us consider the set $\mathrm{Fm}_{\mathbf{IML}}$ of formulas generated by the following grammar:

$$p \mid [B_i = \Gamma|t] \mid [D_i = \varphi|t] \mid [X \leftarrow x|t]\psi \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \psi \to \psi$$

where $p \in Var$, $B_i \in \Upsilon$, $\Gamma \in r(B_i)$, $D_i \in \Delta$, $\varphi \in r(D_i)$ $t \in \mathbb{N}$, $m \geq 1$, $X \in \Delta \cup \Upsilon$ are distinct variables and $x \in r(X)$. Intuitively, the formula $[X \leftarrow x|t]\psi$ means that ψ would hold if X were set to x at time t. In the sequel, formulas of the form $[X_1 \leftarrow x_1|t][X_2 \leftarrow x_2|t] \dots [X_n \leftarrow x_n|t]\psi$ will be denoted by $[X_1 \leftarrow x_1, \dots, X_m \leftarrow x_m|t]\psi$ while, whenever no danger of confusion will be impending, formulas of the form $[X_1 \leftarrow x_1|t_1][X_2 \leftarrow x_2|t_2] \dots [X_n \leftarrow x_n|t_n]\psi$ will be denoted by $[\overrightarrow{X} \leftarrow \overrightarrow{x}|\overrightarrow{t}]\psi$. Formulas of the form $[X = x|t]$ will be called, using a customary terminology, *events*. formulas of the form $[X_1 = x_1|t_1] \wedge \cdots \wedge [X_n = x_n|t_n]$ will be denoted by $[\overrightarrow{X} = \overrightarrow{x}|\overrightarrow{t}]$.

Let $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i \in \mathsf{Ag}}, \{f_i^D\}_{i \in \mathsf{Ag}})$ be an influence model for Ag. $\psi \in \mathrm{Fm}_{\mathbf{IML}}$ is said to be true in $\mathcal{M}$ at time t ($\mathcal{M}, t \models \psi$) according to the following inductive definition.

- $\mathcal{M}, t \models p$ iff $p \in h(t)$, for any $p \in Var$;
- $\mathcal{M}, t \models [B_i = \Gamma|t']$ iff $H_{B_i}^{\mathcal{M}}(t') = \Gamma$;
- $\mathcal{M}, t \models [D_i = \varphi|t']$ iff $H_{D_i}^{\mathcal{M}}(t') = \{\varphi\}$;
- $\mathcal{M}, t \models \neg\psi$ iff $\mathcal{M}, t \not\models \psi$
- $\mathcal{M}, t \models \psi \wedge \chi$ iff $\mathcal{M}, t \models \psi$ and $\mathcal{M}, t \models \chi$;

- $\mathcal{M}, t \models \psi \vee \chi$ iff $\mathcal{M}, t \models \psi$ or $\mathcal{M}, t \models \chi$;
- $\mathcal{M}, t \models \psi \rightarrow \chi$ iff $\mathcal{M}, t \not\models \psi$ or $\mathcal{M}, t \models \chi$;
- $\mathcal{M}, t \models [X_1 \leftarrow x_1, \ldots, X_m \leftarrow x_m | t']\psi$ iff $\mathcal{M}^{[\overrightarrow{X} \leftarrow \overrightarrow{x} | t']}, t \models \psi$.

We observe that the class of influence models over $\mathfrak{S}$ induce a *propositional* consequence relation $\vdash_{\mathbf{IML}} \subseteq \wp(\mathrm{Fm}_{\mathbf{IML}}) \times \mathrm{Fm}_{\mathbf{IML}}$ as follows. Let $\Gamma \cup \{\varphi\} \subseteq \mathrm{Fm}_{\mathbf{IML}}$. We set $\Gamma \vdash_{\mathbf{IML}} \varphi$ iff there is $\Gamma' \subseteq \Gamma$ such that $|\Gamma'| < \omega$ and, for any influence model $\mathcal{M}$, and any $t \in \mathbb{N}$:

$$\text{if } \mathcal{M}, t \models \gamma, \text{ for any } \gamma \in \Gamma', \text{ then } \mathcal{M}, t \models \varphi.$$

As the reader may promptly notice, $\vdash_{\mathbf{IML}}$ is a reflexive, transitive, and monotonic relation (see e.g. [8]).

Lemma 1. *Let $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i \in \mathsf{Ag}}, \{f_i^D\}_{i \in \mathsf{Ag}})$ be an influence model for* Ag. *For any $\varphi \in \mathrm{Fm}_{\mathbf{CL}}$, $t \in \mathbb{N}$ one has that $\mathcal{M}, y \models \varphi$ iff $\varphi \in h(t)$.*

Proof. The proof can be provided by means of a customary induction on the structure of formulas.

In the sequel, our attempt will be to provide a formal treatment of sentences of the form "agent a utters the fake news φ at time t in the influence model $\mathcal{M}$". Our notion will rely on the concept of news. In the present framework, a news is a sentence φ which is relevant to the public debate in the following sense. It *causes* a modification of the audience's beliefs as it is accepted by all or, at least, a part of them, or it is considered as worth being communicated. It is important to observe that our notion of news does not include truth in its definition. A news is a proposition φ that influences the public debate in a social network, in the sense that it influences beliefs and shares of its actors. An informal definition of news can be thus provided as follows. An agent a utters a news φ at time t in the model $\mathcal{M}$ provided that $\mathcal{M}, t \models [D_a = \varphi | t]$ and φ is relevant to a non-empty set A of agents. Namely, in subsequent time steps $t < t_1, \ldots, t_n$ the utterance of φ by a at t causes the acceptance or the spread of φ by agents from A in $t_1, \ldots, t_n$. In other words, φ is news iff it is accepted by agents from A and/or is judged by them to be worth communicating.

We start from the concept of (conditional) influence. As the reader may recognise, it is an adaptation (*mutatis mutandis*) of one of the versions of the concept of *actual causation* outlined by J. Halpern in [10].

Definition 4. *Let $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i \in \mathsf{Ag}}, \{f_i^D\}_{i \in \mathsf{Ag}})$ be an influence model for* Ag, $X_1, \ldots, X_n, Y_1, \ldots, Y_m \in \Delta \cup \Upsilon$, $x_i \in r(X_i)$, $y_i \in r(Y_i)$, $t_1, \ldots, t_n, t''_1, \ldots, t''_m, t \in \mathbb{N}$ *with $t_1 \leq t_2 \leq \cdots \leq t_n < t$ and $t''_1 \leq \cdots \leq t''_m \leq t$ $(n, m \geq 1)$. We say that* $[\overrightarrow{X} = \overrightarrow{x} | \overrightarrow{t}]$ influences *the occurrence of $\varphi \in \mathrm{Fm}_{\mathbf{IML}}$ at t under the condition $[\overrightarrow{Y} = \overrightarrow{y} | \overrightarrow{t''}]$ provided that the following conditions are satisfied:*

1. $\mathcal{M}, t \models [\overrightarrow{Y} \leftarrow \overrightarrow{y} | \overrightarrow{t''}]([\overrightarrow{X} = \overrightarrow{x} | \overrightarrow{t}] \wedge \varphi)$;

2. *There are* $W_1, \dots, W_l \in \Delta \cup \Upsilon$, $t'''_1 \leq \dots \leq t'''_l < t$ *and a setting* $\overrightarrow{x'}$ *of the variables in* $\overrightarrow{X}$ *such that, if* $\mathcal{M}, t \models [\overrightarrow{Y} \leftarrow \overrightarrow{y}|\overrightarrow{t''}][\overrightarrow{W} = \overrightarrow{w}|\overrightarrow{t'''}]$, *then*

$$\mathcal{M}, t \models [\overrightarrow{Y} \leftarrow \overrightarrow{y}|\overrightarrow{t''}][\overrightarrow{X} \leftarrow \overrightarrow{x'}|\overrightarrow{t}][\overrightarrow{W} \leftarrow \overrightarrow{w}|\overrightarrow{t'''}]\neg\varphi;$$

3. $\overrightarrow{X}$ *is minimal, in the sense that there is no proper subset* $\overrightarrow{X'}$ *of* $\overrightarrow{X}$ *such that* $[\overrightarrow{X'} = \overrightarrow{x'}|\overrightarrow{t'}]$ *(where* $\overrightarrow{x'}$ *and* $\overrightarrow{t'}$ *are the restriction of* $\overrightarrow{x}$ *resp.* $\overrightarrow{t}$ *to the variables in* $\overrightarrow{X'}$*) satisfies the first two conditions.*

Specifically, by Definition 4, the notion of conditional influence provides the causal foundation for the characterisation of news: a proposition qualifies as news only if, under certain conditions, it produces actual effects on the declarations or beliefs of other agents. For example, the death of a person is a factual event, but it becomes news only if the person is known by at least someone; otherwise, it remains irrelevant from an informational point of view. Similarly, the discovery of a new particle in a laboratory is an objective fact, yet it becomes news only if it affects the beliefs or actions of scientists, institutions, or the public. In both cases, the underlying mechanism is the same: an event influences a proposition if (a) it actually occurs together with the effect, (b) it is minimal, i.e., not redundant, and (c) keeping fixed some conditions that have indeed taken place, if the event had not occurred, the effect would not have manifested. This provides the background for the next definition, where the notion of news is formally introduced.

Definition 5. *Let* $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i \in \mathsf{Ag}}, \{f_i^D\}_{i \in \mathsf{Ag}})$ *be an influence model for* Ag, $Y_1, \dots, Y_n \in \Delta \cup \Upsilon$ $(n \geq 1)$, $t_1, \dots, t_n, t \in \mathbb{N}$ *with* $t_1 \leq \dots \leq t_n \leq t$, $\varphi \in \mathrm{Fm}_{\mathcal{L}}$ *is a:*

- news *uttered by* $a \in \mathsf{Ag}$ *at* $t \in \mathbb{N}$ *under the condition* $[\overrightarrow{Y} = \overrightarrow{y}|\overrightarrow{t}]$ *provided that the following conditions are satisfied:*
  1. $\mathcal{M}, t \models [D_a = \varphi|t]$
  2. *There are* $a_1, \dots, a_m \in \mathsf{Ag}$ $(m \geq 1)$ *such that* $a_i \neq a$, *for any* $i = 1, 2, \dots, m$, *and there are* $t < t'_1 \leq \dots \leq t'_m$ *such that, for any* $1 \leq i \leq m$, *at least one of the following holds:*
     - $[D_a = \varphi|t]$ *influences* $[B_{a_i} = \Gamma_i|t'_i]$ *under* $[\overrightarrow{Y} = \overrightarrow{y}|\overrightarrow{t}]$, *where* $\varphi \in \Gamma_i \in r(B_{a_i})$ *and* $\varphi \notin H^{\mathcal{M}}_{B_{a_i}}(t-1)$*; or*
     - $[D_a = \varphi|t]$ *influences* $[D_{a_i} = \varphi|t_i]$ *at* t_i *under the condition* $[\overrightarrow{Y} = \overrightarrow{y}|\overrightarrow{t}]$.

     For any $i = 1, 2, \dots, m$, *we call the formula* $[B_{a_i} = \Gamma_i|t_i] \vee [D_{a_i} = \varphi|t_i]$ *and the agent* a_i *a* witness *of the diffusion of* φ.
- news *uttered by* $a \in \mathsf{Ag}$ *at* $t \in \mathbb{N}$ *if it is a news uttered by* a *at* t *under some condition* $[\overrightarrow{Y} = \overrightarrow{y}|\overrightarrow{t'}]$ *such that* $\mathcal{M}, t \models [\overrightarrow{Y} = \overrightarrow{y}|\overrightarrow{t'}]$.

Definition 5 refines the previous notion of conditional influence by introducing the concept of news. A declaration counts as news only when it propagates

beyond the agent who uttered it, generating effects on the beliefs or further declarations of others. The role of witnesses is crucial here: they provide explicit evidence of the diffusion process, showing how an initial statement is causally linked to the informational state of other agents.
Consider, for example, the announcement of a Nobel Prize. The fact that a scientist has been awarded is an event, but it becomes news only if it triggers a cascade of effects: other institutions issue statements, journalists report on it, and members of the public update their beliefs. Each of these reactions constitutes a witness of diffusion, testifying that the information has not remained isolated but has spread through the network.
This definition thus connects the act of uttering a proposition with its social impact: a news is not merely an event or a statement, but a statement whose influence is demonstrated by its ability to reach and affect other agents.
A further remark comes next. As the reader might argue, for a proposition φ to a be a news, Definition 5 requires the existence of witnesses, namely of agents which, in subsequent time steps, update their beliefs according to φ or share φ under the influence of the utterance of φ occurred at a previous step. Therefore, our concept of news seems to rule out, e.g. situations in which an agent, think e.g. of a journalist, is aware of having news at hand before she shares it at a time step. In fact, we argue that being aware of possessing a news item entails anticipating its potential impact on the beliefs and patterns of sharing among its recipients. Consequently, our definition may also be applicable to this case.

In the light of concepts just introduced, one can then provide a definition of fake news as follows.

Definition 6 (Fake news). *$\varphi \in \mathrm{Fm}_{\mathcal{L}}$ is said to be a* fake news *(under conditions $[\overrightarrow{Y} = \overrightarrow{y} | \overrightarrow{t}]$) uttered by a at t provided that:*

F1. φ is a news uttered by a at t (under the condition $[\overrightarrow{Y} = \overrightarrow{y} | \overrightarrow{t}]$) according to Definition 5;
F2. $\mathcal{M}, t' \models [B_a = \Gamma \mid t] \wedge \neg\varphi$ ($\mathcal{M}, t' \models [\overrightarrow{Y} = \overrightarrow{y} | \overrightarrow{t}][B_a = \Gamma \mid t] \wedge \neg\varphi$) with $\varphi \notin \Gamma$, for every $t' \geq t$.

(F1) Rules out the deliberate communication of a piece of information which is completely irrelevant to the public opinion, as e.g. it is rejected by all agents in the network as plainly false, or it is regarded by all as unworthy of dissemination, as it is either uninformative or unreliable.
(F2) Expresses the fact that φ is false (according to the knowledge of the world) and a either knows it is false ($\mathcal{M}, t \models [B_a = \Gamma]$ and $\neg\varphi \in \Gamma$), and in this case, we will talk about of intentional disinformation, or a utters φ without concern for its truth or falsity ($\varphi, \neg\varphi \notin \Gamma$). This last situation covers cases in which false information is spread with goals, e.g. so as to generate economic benefits, different from disinformation.
The next lemma summarises some properties of fake news which are immediate consequences of the definitions provided so far.

Proposition 2. *Let $\mathcal{M} = (\mathfrak{S}, h, \{f_i^B\}_{i \in \mathsf{Ag}}, \{f_i^D\}_{i \in \mathsf{Ag}})$ be an influence model for Ag, $\varphi \in \mathrm{Fm}_{\mathcal{L}}$. The following hold:*

1. *If φ is a fake news uttered by a at t, then it is a news uttered by a at t;*
2. *If φ is a fake news uttered by a at t, then $\mathcal{M}, t \models \neg\varphi$;*
3. *If $\mathcal{M}, t \models [D_a = \varphi | t']$ and, for any $b \in \mathsf{Ag} \smallsetminus \{a\}$, $t'' > t'$, $\mathcal{M}, t \models \neg[D_b = \varphi | t''] \wedge \neg[B_b = \Gamma | t'']$, for any $\Gamma \in r(B_b)$ such that $\varphi \in \Gamma$, then φ is not a news and therefore is not a fake news uttered by a at t';*
4. *If φ is a fake news uttered by a at t whose diffusion is witnessed by $a_1, \ldots, a_n$, and is also witnessed by $a_{n+1}, \ldots, a_l$, then it is also witnessed by $a_1, \ldots, a_n, a_{n+1}, \ldots, a_l$. Then it is also news with the larger set of witnesses $a_1, \ldots, a_n, a_{n+1}, \ldots, a_m$, for every $m \geq n$.* □

It is important to highlight that, in this paper, we do not consider assertions and their implicatures (namely, the way they are understood or the information they convey in the audience's mind) as separate. Therefore, if agent a utters φ at t in order to spread the fake news ψ which a is aware, it will be inferred by other agents in the light of φ, we will just say that a utters ψ. Furthermore, we do not make distinctions between fake news and *partial* fake news, namely news items whose content is partly true and partly false. Thus, propositions such as "J.F. Kennedy was President of the United States and was a reptilian" will be regarded as potential fake news *tout-court*, notwithstanding the fact that they contain, at least in part, true information. Finally, and this is a main concern of our approach, it should be noticed that, unfortunately, Definition 6 allows us to classify as fake news utterances that, indeed, should not. Think, e.g. of pieces of satire. In fact, providing a formal account of the distinction between fake news and satire is a long-standing problem in the literature as it is related to motivations (the presence of a purpose of deception or not) that guide the sharing of information or, as recently shown, the structure of news [5,11]. Since our framework is mostly concerned with investigating fake news from a logical perspective by relying solely on concepts like influence, agents' shares and beliefs, it seems hard to employ it as a tool for discriminating harmless false declarations (aimed at e.g. entertaining users) from attempts of deception. Nevertheless, since the topic calls for a more in-depth reflection, and in order to keep the manuscript within acceptable length limits, we postpone its examination to future work.

3 Conclusion and Future Work

In this work, we have proposed a formal logical framework to analyse the concepts of "news" and "fake news" through the lens of influence within groups of agents, characterised by their beliefs and shares over time. Our approach is grounded in the notion of conditional influence, capturing the causal dependencies between agents' statements and beliefs, and providing a precise definition of news as a statement that produces observable effects on others, manifested through the propagation of beliefs or further declarations. By formalising these mechanisms, our framework offers a conceptual complement to empirical and computational studies on fake news, paving the way for a deeper understanding of information diffusion and its impacts within social networks. The abstract scope of our framework enables a detailed examination not only of the properties of news and fake news, but also of the agents involved in their generation and propagation.

The framework presented here is deliberately abstract, providing a formal foundation for reasoning about the propagation of news and fake news. A natural avenue for future research is to refine and extend this formal apparatus, for instance, by exploring richer logics or alternative notions of influence. In particular, the framework could be expanded to incorporate probabilistic transmission, noisy perception, or agent-based simulations, thereby capturing the inherent uncertainties of information diffusion. Another promising direction is to examine the interaction between news and corrective information, formalising how mechanisms such as fact-checking can modify propagation dynamics. The framework could also be strengthened to keep track of trust dynamics between the agents, weighting the information received based on how trustworthy is the source of information (see, e.g., [14,15] for a general discussion on how trust could be implemented). Finally, the framework could be integrated with simplified computational models to illustrate, at a conceptual level, how different structural conditions shape the spread of fake news.

Acknowledgements

The work of P. Graziani was supported by the Italian Ministry of Education, University and Research through the PRIN 2022 project "Developing Kleene Logics and their Applications" (DeKLA), project code: 2022SM4XC8. The work of D. Fazio and R. Mascella has been funded by the European Union - NextGenerationEU under the Italian Ministry of University and Research (MUR) National Innovation Ecosystem grant ECS00000041 - VITALITY - CUP C43C22000380007.

References

1. Alessandro Aldini. On the modeling and verification of the spread of fake news, algebraically. *Journal of Logic and Computation*, 32(6):1272–1291, September 2022.
2. Alessandro Aldini, Gianluca Curzi, Pierluigi Graziani, and Mirko Tagliaferri. A probabilistic modal logic for context-aware trust based on evidence. *International Journal of Approximate Reasoning*, 169:109167, 2024.
3. Hubert M. Blalock Jr., editor. *Causal Models in the Social Sciences*. Routledge, 2nd edition, 1985.
4. Elena Broda and Jesper Strömbäck. Misinformation, disinformation, and fake news: Lessons from an interdisciplinary, systematic literature review. *Annals of the International Communication Association*, 48(2):139–166, June 2024.
5. Dipto Das and Anthony J. Clark. Satire vs fake news: You can tell by the way they say it. In *2019 First International Conference on Transdisciplinary AI (TransAI)*, pages 22–26, 2019.
6. Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H. Eugene Stanley, and Walter Quattrociocchi. The spreading of misinformation online. *Proceedings of the National Academy of Sciences*, 113(3):554–559, 2016.

7. Valeria Fionda. Logic-based analysis of fake news diffusion on social media. *Social Network Analysis and Mining*, 15(59), 2025.
8. J. M. Font. *Abstract algebraic logic: an introductory textbook*. College Publications, 2016.
9. Sonal Garg and Dilip Kumar Sharma. Linguistic features based framework for automatic fake news detection. *Computers & Industrial Engineering*, 172:108432, 2022.
10. Joseph Y. Halpern. *Actual Causality*. MIT Press, reprint (paperback) edition, February 2016.
11. Jwen Fai Low, Benjamin C.M. Fung, Farkhund Iqbal, and Shih-Chia Huang. Distinguishing between fake news and satire with transformers. *Expert Systems with Applications*, 187, 2022.
12. Alejandro Buitrago López, Javier Pastor-Galindo, and José A. Ruipérez-Valiente. Frameworks, modeling and simulations of misinformation and disinformation: A systematic literature review, 2024.
13. Kai Shu, Deepak Mahudeswaran, Suhang Wang, and Huan Liu. Hierarchical propagation networks for fake news detection: Investigation and exploitation. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 626–637, 2020.
14. Mirko Tagliaferri and Alessandro Aldini. A trust logic for the varieties of trust. In Javier Camara and Martin Steffen, editors, *Software Engineering and Formal Methods*, pages 119–136, Cham, 2020. Springer International Publishing.
15. Mirko Tagliaferri and Alessandro Aldini. From belief to trust: A quantitative framework based on modal logic. *Journal of Logic and Computation*, 32(6):1017–1047, 03 2022.
16. James Woodward. Causal models in the social sciences. In Stephen P. Turner and Mark W. Risjord, editors, *Philosophy of Anthropology and Sociology*, Handbook of the Philosophy of Science, pages 157–210. North-Holland, Amsterdam, 2007.
17. Xinyi Zhou and Reza Zafarani. A survey of fake news: Fundamental theories, detection methods, and opportunities. *ACM Comput. Surv.*, 53(5), September 2020.