THE LINGUISTIC BLIND SPOT OF VALUE-ALIGNED AGENCY, NATURAL AND ARTIFICIAL

TRAVIS LACROIX

ABSTRACT. The value-alignment problem for artificial intelligence (AI) asks how we can ensure that the 'values'—i.e., objective functions—of artificial systems are aligned with the values of humanity. In this paper, I argue that linguistic communication is a necessary condition for robust value alignment. I discuss the consequences that the truth of this claim would have for research programmes that attempt to ensure value alignment for AI systems—or, more loftily, designing robustly beneficial or ethical artificial *agents*.

Keywords — Artificial Intelligence; AI; Value-Alignment Problems; Principal-Agent Problems; Machine Learning; Objective Functions; Normative Theory; Language; Linguistic Communication; Communication Systems; Information Transfer; Coordination; Values; Preferences; Objectives

1. INTRODUCTION

The value-alignment problem for artificial intelligence (AI) asks how we can ensure that the 'values'—i.e., objective functions—of artificial systems are aligned with the values of humanity, writ large (Russell, 2019). One component of this problem is *technical*, focusing on how to properly encode values or principles in artificial agents so that they reliably do what they ought to do—i.e., what we want them to do, or what we *intend* for them to do. Another component of value alignment is *normative*, emphasising what values or principles from normative theory are the 'correct' ones to encode in AI systems (Gabriel, 2020). However, ensuring value alignment for AI systems—or, more loftily, designing robustly beneficial or 'ethical' artificial *agents*—requires more than just translating our best normative theories into a programming language.

I propose and defend the following claim, which I will refer to as 'the **main claim**' throughout this paper.

DEPARTMENT OF PHILOSOPHY, DALHOUSIE UNIVERSITY

E-mail address: tlacroix@dal.ca.

Date: Unpublished draft of July, 2022. Please cite published version, if available.

Main Claim.

Linguistic communication is a necessary condition for robust value alignment.

I discuss the consequences of the **main claim** in the specific context of valuealignment problems *for* artificial intelligence. However, one should note that the **main claim** makes no specific reference to AI systems since I take it to hold in the context of value alignment more generally than just the specific instantiation seen in recent research on safe AI.

In Section 2, I flag some assumptions that I will make in this paper and clarify what is meant by each component of the **main claim**, discussing linguistic communication (2.1), value-alignment problems (2.2), and necessary conditions (2.3) in turn. With this groundwork laid, I provide two arguments in favour of the **main claim** in Section 3. The first (3.1) can be formulated explicitly as follows:

- (1) Principal-agent problems are primarily problems of informational asymmetries.
- (2) Value-alignment problems are structurally equivalent to principal-agent problems.
- (3) **Therefore**, value-alignment problems are primarily problems of informational asymmetries.
- (4) Any problem that is primarily a problem of informational asymmetries requires information-transferring capacities to be solved, **and** the more complex (robust) the informational burden, the more complex (robust) the information-transferring capacity is required.
- (5) **Therefore**, value-alignment problems require information-transferring capacities to be solved, **and** sufficiently complex (robust) value-alignment problems require robust information-transferring capacities to be solved.
- (6) Linguistic communication is a *uniquely* robust information-transferring capacity.
- (7) Therefore, linguistic communication is necessary for sufficiently complex (robust) value alignment.

The second argument (3.2) follows from the failures of the symbolic systems approach to AI in combination with the rigidity of objective functions for aligning values in present-day AI systems. Finally, I discuss some empirical evidence in favour of the **main claim** (3.3). In Section 4, I highlight some logical implications of **main claim**. Importantly, the **main claim** specifies a lower bound on the difficulty of ensuring value alignment for AI. Section 5 concludes.

2. Conceptual Components of the Main Claim

In this section, I provide some basic definitions and clarify some assumptions I will make in this paper. The **main claim** contains three conceptual components *linguistic communication*, (*robust*) value alignment, and the relation between them, *necessity*. These components are treated separately in Sections 2.1, 2.2, and 2.3.

 $\mathbf{2}$

2.1. Linguistic Communication. Before defending the main claim, the first thing to clarify is what is meant by 'linguistic communication'. Linguistic communication, in this context, means something like 'natural language'. That is to say, 'linguistic communication' is understood as synonymous with 'natural language', and I may use the two terms interchangeably throughout this paper.¹ Note that no *particular* linguistic communication system—e.g., English—will be necessary for value alignment in the context of the main claim; instead, 'linguistic communication' and 'natural language' should be understood as a shorthand for the *abilities* that give rise to linguistic communication—i.e., the *capacity* for language. The capacities that give rise to linguistic communication' ('simple communicative abilities'), typ-ified by the ubiquitous signalling systems in the natural world, from bacteria to bees to primates.²

At the same time, it is perhaps too strong to require that all features of natural language are individually necessary for robust value alignment between agents. Rather, certain constitutive (and differentiating) features of linguistic communication will be necessary for robust value alignment. I will discuss more precisely which functions of linguistic communication are necessary preconditions below. So, not every feature of linguistic communication will be necessary for robust value alignment; however, those features that *are* necessary for robust value alignment are also the features of linguistic communication that make it (many researchers suppose) unique to humans.

A few things are presupposed, which I will clarify but not defend at length. First, I assume that the primary *use* of natural language is *to communicate*. This view is distinct from the views of, e.g., Chomsky (1965, 1980, 1995), who suggests that the primary purpose of language is the expression of thought.³ Those who think that language is not (primarily) for communication argue that it appears that only *some* aspects of language could be understood as adaptations *for* communication (Fitch, 2010, 21). Thus, for Chomsky and those who follow the biolinguistic research programme—i.e., mainstream generative grammar—human language is identified as something like the FLN—the *faculty of language in the*

¹Of course, it is logically possible that there is a form of linguistic communication that is *not* a natural language, thus pulling these two concepts apart. However, natural language is a *pro-totypical* example of linguistic communication. As a result, in the context of this paper, we can understand them as relative synonyms.

²See, e.g., Lishak (1984); Lugli et al. (2003); Marler and Slabbekoorn (2004); Belanger and Corkum (2009); Houck (2009); Mäthger et al. (2009); Bruschini et al. (2010); Costa-Leonardo and Haifig (2010); Haddock et al. (2010); Wyatt (2010); Thiel and Breithaupt (2011).

³When Chomsky discusses 'language' (or, more typically, 'grammar'), he means '*I*-language'—i.e., a bio-computational system represented in the brain, with a capacity for generating a discrete infinity of hierarchical structures. See also, Bickerton (1990); Wray (1998); Hauser et al. (2002).

narrow sense (Hauser et al., 2002). The FLN emphasises the computational mechanisms required for (complex or compositional) syntax.⁴ But, it has been argued that this 'syntactocentrism'—where syntax is supposed to be the only generative component of language—is misplaced.⁵ There is no obvious conceptual clarity to be gained by assuming that language is different *in kind* from communication. Furthermore, assuming that communication is the primary purpose of language is more parsimonious (in an evolutionary context) than alternative views, which assume it is not.⁶ The view that language is primarily communicative further underscores the *continuity* between human linguistic abilities and the non-linguistic communicative abilities of non-human animals.

This last point is relevant to discussions of value alignment *qua* moral behaviour insofar as an analogy holds between language and ethics. Briefly, linguistic ability and moral behaviour are typically considered unique to humans.⁷ However, some non-human animals exhibit complex communication abilities, which contain features that might be called *proto-linguistic*.⁸ Additionally, some non-human animals exhibit *proto-moral* behaviours in the form of the complex social structures, norms, and reactive emotions that give rise to robust cooperation (of the sort theorised by some as providing a foundation out of which human morality may have evolved).⁹

It is difficult to specify precisely what differentiates human-level linguistic communication (abilities) from simpler communication systems seen in non-human

⁴Hauser et al. (2002) hypothesise that the FLN is equivalent to recursion, which makes the FLN unique to *Homo sapiens*. This idea is sometimes referred to as the 'recursion only' hypothesis i.e., the hypothesis that recursion is *the* property that 'distinguishes human language from animal communication systems' (Traxler et al., 2012, 611). See criticism in Jackendoff and Pinker (2005); Pinker and Jackendoff (2005); Parker (2006); Wacewicz (2012).

 $^{^{5}}$ See discussion in Jackendoff (1997, 2003, 2012); Culicover and Jackendoff (2005). When the FLN is introduced by Hauser et al. (2002), its uniqueness to *Homo sapiens* is postulated as a hypothesis to be bolstered with empirical data. However, Fitch et al. (2005) later *stipulate* that the FLN is unique to humans as a matter of definition. Wacewicz (2012) highlights that the description of the FLN proffered by Fitch et al. (2005) is not merely a correction or clarification of the earlier definition given by Hauser et al. (2002); instead, they are conceptually inconsistent. Thus, the distinction between the *faculty of language in the broad sense* (FLB) and the FLN confuses, rather than clarifies, dialogue on language origins. Fitch (2017), in particular, appears to have somewhat distanced himself from his early work with Hauser and Chomsky, referring instead to the *derived components of language* (DCLs), which include (at least) complex vocal learning, hierarchical syntax, and complex semantics/pragmatics. Thus, Fitch (2017) simultaneously de-emphasises the importance of syntax alone.

⁶See discussion in LaCroix (2020).

⁷For example, Kitcher (2011); Korsgaard (2018) argue that humans are the only animals with the sort of meta-cognitive capacities required for moral agency.

⁸See discussion in Bickerton (1990, 1995); Wray (1998, 2000, 2002); Jackendoff (1999, 2003); Arbib (2002, 2003, 2005); Tallerman (2007, 2012); Planer and Sterelny (2021).

⁹See discussion in Kitcher (2006a,b, 2011); de Boer (2011) and empirical examples given by de Waal (1996, 2006); Flack and de Waal (2000). Vincent et al. (2019) argue that great apes participate in certain (foundational) normative practices like reciprocity, caring, social responsibility, and solidarity; similarly, Rowlands (2012) argues that some non-human animals can track moral truths. Of course, this debate may be semantic rather than substantive, as Fitzpatrick (2017) argues.

species. Nonetheless, certain distinctive features (or functions) of language, understood as a highly sophisticated communication system, will be necessary for the robust sort of value-alignment to which the **main claim** refers. One such feature of linguistic communication is *compositionality*—and related features like hierarchy and recursion—which is often taken as a key differentiating component of linguistic communication systems insofar as it is absent in any hitherto studied animal communication system.¹⁰ 'Compositionality', here, can be taken as the usual definition, which is given by the following principle:

Principle of (Linguistic) Compositionality.

The meaning of a compound [complex] expression is a function of the meaning of its parts [constituents] and how they are combined [composed].¹¹

The key thing here is that the elements of natural language can be combined into hierarchical phrases, which then may be recursively combined into larger phrasal expressions; and, the meaning of such an expression is a *function* of the meaning of its parts and how they are combined. The relation between recursion, hierarchy, and compositionality is such that recursion requires hierarchy—at least to some extent—and hierarchy requires compositionality—again, at least to some extent (De Beule, 2008). A key entailment of a communication system's being compositional is that, with a limited vocabulary and a finite set of grammatical rules, linguistic communication systems allow for the production and understanding of an unlimited number of unique expressions.

The principle of (linguistic) compositionality contrasts with a notion of *trivial* compositionality, where one can always interpret the complex expressions of a communication system in terms of set intersection.¹² For example, the words 'cute' and 'dog' can compose to form the expression 'cute dog' in English; however, the meaning of this expression is determined trivially by the intersection of the sets

¹⁰See, e.g., discussion in Hauser et al. (2002); Hauser and Fitch (2003); Mehler et al. (2006); Fitch (2010); Hurford (2012); Scott-Phillips and Blythe (2013); Berwick and Chomsky (2016). Of course, some non-human animals utilise remarkably complex communication systems—e.g., the combinatorial 'waggle dance' of *Apis mellifera* (Aristotle, 1995; Spitzner, 1788; von Frisch, 1967); the syntactic songs of *Poecile atricapillus* (Hailman et al., 1985); the 'pyow-hack' signals of *Cercopithecus nictitans* (Arnold and Zuberbühler, 2006a,b, 2008, 2013); or the functionally-referential alarm calls of *Chlorocebus pygerythrus* (Seyfarth et al., 1980b) and *Cercopithecus diana* (Zuberbühler et al., 1999), among others. Still, despite this remarkable variation and complexity, non-trivial (linguistic) compositionality appears to be unique to human-level linguistic communication systems.

¹¹See Frege (1923); Partee (1984); Kamp and Partee (1995); Janssen (1997, 2012); Pagin and Westerståhl (2010a,b); Szabó (2012, 2020). Of course, this formulation is inherently ambiguous: Szabó (2012) highlights that these ambiguities give rise to at least eight different readings of the principle.

¹²See Schlenker et al. (2016); Zuberbühler (2018); Steinert-Threlkeld (2020).

of *cute-things* and *dog-things*.¹³ The principle of (linguistic) compositionality also contrasts with the individual notions of *syntactic* and *semantic* compositionality. A complex signal is syntactically compositional just in case it is composed of atomic signals but is not interpreted compositionally. Conversely, it is semantically compositional just in case the meaning of the compositional signal can be *decomposed* in interpretation, even if the signal is holophrastic. Linguistic compositionality requires both semantic and syntactic composition.¹⁴

Compositionality is typically understood as a requirement for explaining both the *systematicity* and *productivity* of linguistic communication.¹⁵ In any case, I will take it for granted that linguistic communication systems, typified by natural language, are compositional in a non-trivial way—i.e., the robust form of compositionality that appears to be a key feature of linguistic communication involves some *non-conjunctive*—or what we might call *functional*—modification of linguistic items by certain other linguistic items, such as function words (Steinert-Threlkeld, 2020). Composition, hierarchy, and recursion are sufficient for systematicity and productivity; so, compositionality, for the claims made in this paper, will be taken as a defining feature of linguistic communication.¹⁶

To summarise, by 'linguistic communication', I mean the particular *types* of communication systems that evolved in *Homo sapiens*. These species-unique communication systems involve (verbal or nonverbal) signals and their conventional meanings; they give rise to hierarchical, recursive, compositional structures; and, their primary function is to allow individuals to share information—i.e., to communicate. In short, linguistic communication systems are those communication systems that are typified by natural languages. I will discuss the importance of the non-trivially compositional features of linguistic communication systems for the **main claim** in Section **3**. First, I describe what is meant by 'robust value alignment' in the **main claim**.

2.2. Robust Value Alignment. The next thing to clarify is what is meant by 'robust value alignment' in the context of the main claim. I begin by discussing the *value-alignment problem* in its most general form, highlighting a range of conditions under which tokens of this type of problem may be generated (2.2.1). I then turn to the value-alignment problem in the particular context of AI (2.2.2). Finally, I discuss the connection between the value-alignment problem and the *control* problem,

 $^{^{13}}$ See **here** for an exemplar.

¹⁴See discussion in LaCroix (2020, 2022a).

¹⁵See discussion in Fodor (1998); Szabó (2020).

¹⁶Note that this can be true regardless of whether the thing that evolved in *Homo sapiens* was compositionality per se. LaCroix (2020, 2021, 2022c) argues that *reflexivity* is a more apt explanatory target for a (gradualist) evolutionary account of language origins. Reflexivity can give rise to compositionality as a byproduct.

arising in the context of superintelligent AI systems (2.2.3). As in Section 2.1, the main goal of this section is to clarify some basic concepts, definitions, and assumptions in the context of this paper. I will again postpone discussing the importance to the **main claim** of the features of value-alignment problems discussed in this section to Section 3.

2.2.1. General Value-Alignment Problems. The value-alignment problem, in its most general form, is a problem of how two (or more) agents (actors) can align their values (or objectives). In economics, law, and politics, what I am calling the value-alignment problem is more commonly known as the *principal-agent problem*.¹⁷ This type of problem arises (or, at least, may arise) in any context where some entity—called 'the principal'—appoints another entity—called 'the agent'—to act on the principal's behalf. A problem may be generated, in part, because the principal and the agent have different objectives, incentives, or values. When there is a conflict of interest between the principal and the agent's values, we say their values are misaligned—hence this is a value-alignment problem. Thus, in its most general form, the *problem* of value alignment arises from the dynamics of multi-agent interactions involving the delegation of tasks from one actor to another.

Some examples will clarify how ubiquitous this problem is in human-human interactions.

- (1) A corporate executive, like a chief executive officer (CEO), runs a corporation on behalf of its shareholders. The shareholders are the principal(s), and the CEO is the agent.
- (2) In a democracy, citizens elect a representative (or representatives) who then act on behalf of the citizens. The citizens are the principal(s), and the elected party or officials are the agent(s).
- (3) A customer orders groceries online, and a grocer collects the items to be delivered. The customer is the principal, and the grocer is the agent.
- (4) An individual's car performs sub-optimally, so they take it to a mechanic, who tells them several parts need to be replaced. The car owner is the principal, and the mechanic is the agent.
- (5) A homeowner uses a real estate agent to sell a house. The homeowner is the principal, and the real estate agent is the agent.

In case (1), we might imagine that the objective of the shareholders is for the corporation to maximise (shareholder) profit through increasing stock value. However,

¹⁷This problem is sometimes referred to as an *agency dilemma* or an *incentive problem*; see discussion in Jensen and Meckling (1976); Eisenhardt (1989); Laffont and Martimort (2002).

the CEO might use profits to proffer large bonuses to corporate-level executives instead of paying dividends to shareholders—an action that benefits the agent(s) but not the principal(s). In case (2), citizens might vote for a party whose platform (appears to) align with their values; however, once elected, the party may renege those promises in light of competing considerations that benefit the party rather than the citizens. In case (3), we might imagine that a customer values certain properties in their groceries—e.g., freshness—whereas the grocer values certain other things—e.g., getting rid of items close to expiration. In case (4), the car-owner may value having all and only the required work performed, and the mechanic may value maximising unnecessary expense to their own benefit. In case (5), the homeowner may value getting the best price on the house, whereas the agent may value closing the sale as quickly as possible (even at a reduced price); or, the agent may overpromise on the price to win the listing in the first place, simultaneously making it more difficult to sell the property.

A principal-agent problem can arise in each case because the principal and the agent may have different values, interests, or objectives. However, another key feature common to these examples is *informational asymmetry*.¹⁸ In case (1), the shareholders do not have information about the day-to-day goings-on of the business, but the CEO does. In case (2), the citizens use platforms as a proxy for choosing the candidate whose values most align with their own while not knowing which values the party will instantiate once elected. In case (3), the principal cannot observe the agent's actions. In cases (4) and (5), the principal lacks specialised information, which allows the agent to take advantage for their own gain. Informational asymmetries and imperfect information contribute to the generation of principal-agent problems.¹⁹

Laffont and Martimort (2002) highlight that if there is no private information between a principal and an agent, then *even if* the agent's objectives conflict with the principal's, the principal could still propose 'a contract which perfectly controls the agent and induces the [agent's] actions to be what he [the principal] would like to do himself in a world without delegation' (12). Essentially, under complete information, the principal has complete *bargaining power*.²⁰ Therefore, misaligned

 $^{^{18}}$ In the 1970s, economists showed how asymmetric information poses significant challenges to General Equilibrium Theory (Akerlof, 1970; Spence, 1974; Rothschild and Stiglitz, 1976). See also Marschak and Radner (1972) and discussion in Laffont and Martimort (2002).

¹⁹Note that 'perfect information' is a term of art. In Economics, perfect information implies that all market participants have all the information required to make a decision. In game theory, perfect information means that a player knows the game's entire history up to the decision point, as in backgammon. Imperfect information is the negation of perfect information; this occurs when some information is unavailable or hidden. Thus, imperfect or incomplete information means that there is some uncertainty. See discussion in von Neumann and Morgenstern (1944); Shapley (1953). ²⁰In classical game theory, this is often operationalised as a higher disagreement point for the powerful agent; see discussion in LaCroix and O'Connor (2021).

values alone are insufficient for generating a principal-agent problem since they can be controlled when there is no informational asymmetry between the principal and the agent.²¹ There are three different ways that private information can generate an agency dilemma.

First, *hidden knowledge* is an informational asymmetry resulting from the agent's private information—e.g., concerning their own skills or opportunity costs—to which the principal does not have access. Hidden knowledge may contribute to the generation of a principal-agent problem.²² For example, a prospective employee (the agent) knows their background skill level and appropriateness for a particular job, whereas the hiring committee (the principal) does not. In cases of hidden knowledge, uncertainty is exogenous to the relationship between the principal and the agent. Cases (4) and (5) above are examples of informational asymmetries arising from hidden knowledge.

Second, *hidden action* is an informational asymmetry caused by the agent's ability to perform an action that the principal cannot observe.²³ When the risk-taking individual (the agent) knows more about their intentions than the consequencepaying individual (the principal), the agent may take on more risk than the principal would otherwise be comfortable with, as is common with insurance. In this case, the uncertainty due to asymmetric information is endogenous to the relationship of the principal and the agent. Cases (1) to (3) above exemplify informational asymmetries arising from covert action.

Solutions to principal-agent problems arising from informational asymmetries due to hidden action and hidden knowledge assume that information is (ex post) verifiable by an independent third party, such as a (benevolent) Court of Justice (Laffont, 2000). However, a third type of informational asymmetry may give rise to an agency dilemma when we assume that the information between an agent and principal is symmetric (ex post) but unverifiable *in principle* by a third party (Sappington, 1991). Thus, the third type of informational asymmetry that may generate a principal-agent problem arises from the non-verifiability of (otherwise symmetric) information.²⁴ For example, we might suppose that the principal and the agent have

²¹This claim is mathematically provable on the particular economic model we are discussing. Of course, since this is a model and, therefore, an idealisation, we might question whether this model is sufficiently applicable to real-world interactions and whether this claim holds in the real world. ²²In economics, this is referred to as *adverse selection*. Hidden knowledge on the agent's part causes the principal to give up some information rent. Thus, a contract must be designed to elicit private information, which may be costly to the principal. See Akerlof (1970); Rothschild and Stiglitz (1976); Spence (1973, 1974); Laffont and Martimort (2002); Hou et al. (2009).

 $^{^{23}}$ In economics, this is known as *moral hazard*. See Haynes (1895); Knight (1921); Arrow (1963, 1968); Vaughan (1997); Laffont and Martimort (2002).

 $^{^{24}}$ Non-verifiability is particularly relevant in the field of *contract law*, though Shah (2014) notes that non-verifiability receives much less coverage in the economic literature on principal-agent problems than hidden knowledge (adverse selection) and hidden action (moral hazard). See

identical and sufficient (ex ante) information to complete some transaction. Here, a principal-agent problem may still arise when the agent represents the true state of the world as being different than they (and the principal) know it to be. This is a problem when either the agent's representation of the world is unverifiable by a third party *or* when it is too costly for a third party to verify.

Thus, value misalignment may give rise to a principal-agent problem. However, when information is symmetric, the principal can create a contract that induces the agent to act just as the principal would in the absence of delegation. Therefore, value misalignment *alone* is not sufficient to generate a principal-agent problem.

Furthermore, I suggest that value misalignment is also *unnecessary* to generate a principal-agent problem. For example, suppose the agent and principal have perfectly-aligned objectives, but they have no way to transmit that information. In that case, it is still possible that the agent's actions misalign with the principal's objectives (despite the agent's intentions). This situation might occur if, for example, there is an optimal action which would satisfy the principle's objectives—and, *ex hypothesi*, would also satisfy the agent's objectives—but the existence of this action is not common knowledge.²⁵ Simple coordination games, where the agent chooses an action and the principal and agent both receive the same payoff, provide a clear example of this possibility.²⁶

To summarise, a principal-agent *problem* may arise in any situation where an entity (the agent) can make decisions or take actions on behalf of, or that impact, another entity (the principal). The delegation of tasks from the principal to the agent generates a problem of managing information flows. Therefore, principal-agent problems are fundamentally problems of asymmetric information rather than asymmetric values (though value misalignment may exacerbate these problems).²⁷

Williamson (1973, 1975); Grossman and Hart (1986); Sappington (1991); Hart (1995); Laffont (2000); Laffont and Martimort (2002).

 $^{^{25}}$ Suppose the philosophical literature on peer disagreement is to be trusted. In that case, it is at least possible for this to be true even when both parties are privy to identical evidence.

²⁶As a toy case, suppose the principal chooses a number, and the agent has to guess the correct number for both players to receive payoff. Their values are perfectly aligned, but the agent may still act in a way that the principal would not, precisely because there is an informational asymmetry between them. (With thanks to Aydin Mohseni for raising this possibility to me.) This toy example is much less artificial than it may seem at first glance—simple signalling games have a similar structure.

²⁷The preceding discussion should make clear that I am using certain terms as relative synonyms for example, 'objectives', 'values', and 'interests'. Although there may be relevant conceptual distinctions between these terms, I will ignore these for ease of exposition in this paper. When I discuss artificial agents, I will refer to 'their values' without worrying too much about anthropomorphism; I will take the 'values' of an AI system to mean something like the 'objectives'—in the sense of *objective functions*—of said system.

2.2.2. Value Alignment Problems for AI. In the preceding discussion of principalagent problems from economics, each example involved interactions between a human principal and a human agent. However, this class of problems also arises in the context of artificial agents or AI systems, where it is typically referred to as the 'value-alignment problem'. Hadfield-Menell and Hadfield (2019, 417) suggest that the value-alignment problem has a 'clear analogue' in principal-agent problems. However, we can make a stronger claim: the relationship between value-alignment problems for AI and principal-agent problems in an economic context is not an analogy but identity—at least in terms of formal structure. However, the structural identity between these problems does not imply that their solutions will be identical, as we will see.

For clarity, in the remainder of this paper, I will use 'the principal-agent problem' to refer to the class of problems arising from conflicts of interest and asymmetric information in general interactions between two entities, as discussed in Section 2.2.1. I will use 'the value-alignment problem' to refer to the subset of principal-agent problems arising in human-AI interactions.

'Artificial intelligence' refers to a property of an artificial system—i.e., that it 'thinks', 'acts', or 'behaves' in an intelligent way. Intelligence, in this context, is best understood as 'an agent's ability to achieve goals in a wide range of environments' (Legg and Hutter, 2007, 12).²⁸ 'AI' also refers to an approach or set of techniques for *achieving* this property in an artificial system (Gabriel, 2020). The most promising method for achieving machine intelligence in recent years is machine learning (ML). This approach to AI involves training models using (typically huge amounts of) data. The models 'learn' gradually to behave in the desired way without that behaviour being explicitly programmed. The three main techniques for ML today are supervised learning, unsupervised learning, and reinforcement learning (Goodfellow et al., 2016).

In supervised learning, a model is trained on labelled examples—the *training* data. The model is evaluated on how well it generalises what it learns from the training data to previously unseen examples—the *test data*. In unsupervised learning, an algorithm learns underlying patterns or correlations from unlabelled data. Reinforcement learning (RL) depends upon sparse rewards for actions (Sutton and Barto, 2018). Since 2012, the main driver of AI research has been *deep learning* (DL). This approach to ML utilises deep neural networks modelled (roughly) after neurons in the human brain (Savage, 2019). DL uses layers of algorithms to process data—information is passed through each subsequent layer in a neural network, with the previous layer's output providing input for the following layer. One of the

²⁸See also Gardner (2011); Cave (2017); Gabriel (2020); Russell and Norvig (2021).

key advantages of DL techniques is that they do not require the heavily hand-crafted features used by traditional methods for AI (Buckner, 2019).

ML, at its core, is just an *optimisation problem*. The thing being optimised, in this case, is an *objective function*, which gives a way to configure the system to bring it closer to the ground truth provided during model training. Essentially, an objective function provides a proxy for what we want the system to do. Consider an example from supervised learning for image recognition. Given a dataset, **D**, of image-label pairs, (x, y), the true objective—i.e., what we want or intend for the system to do—is to have the model correctly guess the image label for previously unseen images. The objective function, in this case, might be a probability distribution function,

$$p(\hat{y} = y \,|\, x, \theta),$$

which outputs the conditional probability that the predicted label, \hat{y} , is identical to the true label, y, given the observed data, x, and a model, θ .

The true label, y, determines whether a particular predicted label, \hat{y} , is correct for a particular image, x. A loss function is used to optimise the model, θ . According to a specified evaluation metric—e.g., mean-squared error²⁹—the loss function tells us how close the model is to the correct prediction for a single point.³⁰ Thus the *true* objective—correctly labelling images—is approximated by an objective function, which determines a metric for how close the model is to the objective.³¹ If loss is close to zero, then the model successfully optimises the objective function according to the metric used. Further, if the objective function is a good proxy for the true objective, then optimising an objective function means optimising the objective. Therefore, a given system's objective function(s) must be accurately specified (Reed and Marks, II, 1999).

In the context of AI, the value-alignment problem can arise when objectives are misspecified. Part of the difficulty here arises because objectives require programmers to define an objective function, which can be difficult to operationalise in a programming language. Additional complications arise from the fact that objective functions are mere proxies for the true objective; however, the objective function

²⁹The mean squared error is given by $\frac{1}{n} \sum_{i=1}^{n} (Y_i - \hat{Y}_i)^2$, where *n* is the number of data points, Y_i is the set of observed values, and \hat{Y}_i is the set of predicted values.

 $^{^{30}}$ The loss function is sometimes differentiated from a *cost function*, which tells us the *average* loss over the entire dataset. This is sometimes couched in the language of *empirical risk*, which describes the loss over all samples in the dataset, which is then contrasted with the *true risk*—i.e., the loss over all samples. More accurately, an objective function approximates our true objective, and the objective function determines the empirical risk, which approximates the true risk. See discussion in Goodfellow et al. (2016, Sec. 8.1).

 $^{^{31}}$ In fact, this *approximates* an approximation. (With thanks to Michael Noukhovitch for bringing this to my attention). Loss is measured on the test set, which gives a measure of generalisation error for a given objective function. Hence, optimisation actually occurs on the training set. Still, the hope is that we are also optimising the test set—i.e., achieving some degree of generalisation. See discussion in Goodfellow et al. (2016, Sec. 5.2).

is identical to the true objective from the AI system's 'point of view'. When objective functions are poorly specified, this will lead to a value-alignment problem insofar as there is a misalignment between the actual objective—our 'values'—and its proxy—the system's 'values'.

The provision of an objective function implicitly defines an optimisation 'landscape' (Sipper et al., 2018); however, the solution space for an optimisation problem defined by an objective function includes a 'pathology' of *local* optima (Lehman and Stanley, 2008), meaning that these landscapes are frequently 'deceptive' (Goldberg, 1987; Mitchell et al., 1992). Lehman and Stanley (2008) highlight that the objective function 'does not necessarily reward the stepping stones in the search space that ultimately lead to the objective' (329), meaning that objective functions are often constructed *ad hoc*. Thus, poorly-designed objective functions can lead to valuealignment problems whenever there is a conflict between the objective function and the actual objective. Furthermore, even when objective functions are accurate proxies, sufficiently complex action spaces will have local optima that may result in outputs that are grossly misaligned with the true objective. A mismatch between a well-specified objective function and the emergent behaviour of an AI system is sometimes referred to as 'inner alignment'; when the objective is misspecified, this is referred to as 'outer alignment' (Hubinger et al., 2021).³²

These value-alignment problems fall under the heading of 'AI safety'. Amodei et al. (2020) highlight that when objective functions are misspecified, this may give rise to unanticipated side effects or reward hacking; when objective functions are too expensive to evaluate at regular intervals, this creates a problem of scalable supervision; and, local optima of objective functions may lead to undesirable behaviour during learning.³³ These problems are exacerbated because the utilities determined by any concretely-specified objective function will necessarily be a mere *subset* of our utilities—i.e., the things we value (Yudkowsky, 2018). This difficulty leads to (what Russell (2019, p. 9) sees as) the failure of the 'standard model' of intelligence for AI systems. Namely, a standard definition of intelligence in humans might be formulated as follows:

Humans are **intelligent** to the extent that **our** actions can be expected to achieve **our** objectives.

Given this definition, the 'standard model' for *machine* intelligence has been given analogously, thus:

³²See also discussion in Ecoffet et al. (2020); Christian (2020); Krakovna et al. (2021); Amodei and Clark (2016).

 $^{^{33}}$ Amodei et al. (2020) refer to these problems as 'accidents', but how they can be classed as value-alignment problems should be clear. See discussion in Hadfield-Menell and Hadfield (2019).

Machines are **intelligent** to the extent that **their** actions can be expected to achieve **their** objectives.

The problem is that, unlike humans, machines have no objectives of their own. Thus, we must define their objectives, which gives rise to the possibility of value misalignment.

Each of the problems described above involves encoding objectives in an AI system. Gabriel (2020) refers to this as the 'technical component' of the valuealignment problem—namely, how do we encode values, principles, objectives, etc. in AI systems so that they 'reliably do what they ought to do' (412), or what we *intend* for them to do. Gabriel (2020) further distinguishes the technical component of the value-alignment problem from the 'normative component', which involves the problem of determining *what* values (objectives) should be encoded in an AI system in the first place.³⁴

Part of the idea is that as these systems become more integrated into society, some of their decisions may carry moral weight so that we might classify their actions as 'moral' or 'immoral'. These considerations have given rise to the field of *machine ethics*, which seeks to 'implement moral decision-making faculties in computers and robots' (Allen et al., 2006).³⁵

Value-alignment problems are truly ubiquitous. 'Accidents', like those described in Amodei et al. (2020), are (technical, internal) value-alignment problems to the extent that misspecified or costly objective functions may lead to behaviour misaligned with what we intended the system to do. More generally, machine bias and problems arising from fairness considerations are value-alignment problems at least to the extent that we do not want or intend for models to act in ways that we would call 'racist', 'sexist', or otherwise discriminatory.³⁶

Note that all the problems mentioned have the same structure as the more general *principal-agent problems* described in Section 2.2.1. Although misaligned values (objectives) can exacerbate these problems, the real cause of a value-alignment problem is asymmetric information. Complex computational systems, like those that undergird AI systems, are often 'ineliminably opaque' (Creel, 2020, 568). Opacity implies that verifying alignment for sufficiently complex systems or situations may

³⁴It is worth noting that Gabriel (2020) uses the language of 'artificial agent' rather than 'AI system' when discussing value alignment. A 'standard' philosophical account of agency requires something like *intentional action*. This view is defended by, e.g., Davidson (1963, 1971); Goldman (1970); Brand (1984); Bratman (1987); Dretske (1988); Bishop (1989); Mele (1992, 2003); Enç (2003), among others. See discussion in Schlosser (2019). However, even a more inclusive notion of agency is unnecessary for generating value-alignment problems. As these systems are further integrated into society, this problem becomes more pressing (LaCroix and Bengio, 2019).

 $^{^{35}}$ See Tolmeijer et al. (2020) for a recent survey of machine ethics and approaches to 'artificial moral agency' (AMA).

³⁶Some salient examples are described in, e.g., Angwin et al. (2016); Christian (2020); Tomasev et al. (2021); Miceli et al. (2022).

be impossible. This difficulty is worsened when there is no simple matter of fact about what the system should align with in the first place.³⁷ These considerations echo the insight from economics that 'by definition the agent has been selected for his specialized knowledge and the principal can never hope to completely check the agent's performance' (Arrow, 1968, 538). Therefore, *transparency, interpretability*, and *explainability* (Creel, 2020, 2021; Erasmus et al., 2021; Kasirzadeh, 2021) are also fundamentally problems of value alignment. These considerations further underscore the insights discussed in Section 2.2.1 that *informational asymmetries* generate principal-agent problems.

Essentially, *any* problem endogenous to the model can be considered a valuealignment problem.³⁸ Even problems of *control*, which arise primarily in the context of superintelligent AI, are a type of value-alignment problem insofar as a lack of control of a system that is capable of modifying its own reward function implies a misalignment of values (again, to the extent that we value maintaining control over such a system).³⁹ As it turns out, the problem of value alignment and the problem of control are closely connected.

2.2.3. The Connection Between Value Alignment and Control. Bostrom (2014) asks the following question: 'how can the sponsor of a project that aims to develop superintelligence ensure that the project, if successful, produces a superintelligence that would realize the sponsor's goals?' (55). This is clearly a question of value alignment, where the principal is a project sponsor, and the agent is a superintelligence. Bostrom (2014) acknowledges the connection between this particular valuealignment problem and principal-agent problems from economics. The problem can be divided into two components. The first, which Bostrom (2014) calls 'the first principal-agent problem', is a principal-agent problem centred on human-human interactions, where the principal is the sponsor of a research project aimed at creating a superintelligence and the agent is the developer of that project. Bostrom (2014) notes that the sponsor (the principal) could be a single individual or the entirety of humanity. The important thing is that the sponsor delegates authority to the developer, giving rise to a principal-agent problem.

³⁷See discussion in LaCroix (2022b); LaCroix and Luccioni (2022).

 $^{^{38}}$ Considering the wider class of interactions that arise from principal-agent problems, negative consequences exogenous to the model might also be called value-alignment problems. However, these will arise primarily from interactions between users of these systems and the stakeholders affected. The prototypical example is *misuse*, where a bad-faith actor intentionally uses technology for harm. However, endogenous value misalignment can be more subtle—as when a research group incentivised by, e.g., profit creates a system that adversely affects some subset of society. See discussion in Falbo and LaCroix (2021) for some examples.

³⁹See discussion in Bostrom (2014); Soares et al. (2015); Orseau and Armstrong (2016); Hadfield-Menell and Hadfield (2019); Russell (2019).

However, a different problem arises from the interaction between the *project* and the *system*—i.e., the superintelligence the project creates. This is the *control problem*, which Bostrom (2014) also refers to as 'the second principal-agent problem'. Effectively, the principal must ensure that the intelligent system it creates will not act in ways that harm the principal's interests. Bostrom (2014) highlights that 'standard management techniques'—i.e., those discussed in Section 2.2.1—apply to the first principal-agent problem. In contrast, managing the second principal-agent problem will require brand-new techniques.⁴⁰

Whereas the first-principal agent problem is ubiquitous, the second principalagent problem arises only in the context of a *superintelligent* AI system. Control is a problem because several instrumental goals—e.g., self-preservation, goal-content integrity, cognitive enhancement, resource acquisition, etc.—are useful sub-goals for almost any original objective we may have programmed in the system (Russell, 2019, 141).⁴¹ Thus, one need not explicitly program incentives for self-preservation in an AI system.

However, the control and value-alignment problems are closely related, as Wiener (1960) highlights:

If we use, to achieve our purposes, a mechanical agency with whose operation we cannot efficiently interfere once we have started it, because the action is so fast and irrevocable that we have not the data to intervene before the action is complete [i.e., the control problem], then we had better be quite sure that the purpose put into the machine is the purpose which we really desire and not merely a colorful imitation of it [i.e., the value-alignment problem]. (1358)

Essentially, if we cannot solve the control problem before the advent of a sufficiently powerful AI system, we had better have solved the value-alignment problem before this happens so that the system's uncontrolled (or uncontrollable) actions at least align with our values.⁴²

16

 $^{^{40}}$ Some such techniques are discussed in Chapter 9 of Bostrom (2014).

⁴¹Bostrom (2014) refers to this as the *instrumental convergence thesis*:

The Instrumental Convergence Thesis

Several instrumental values can be identified which are convergent in the sense that their attainment would increase the chances of the agent's goal being realized for a wide range of final goals and a wide range of situations, implying that these instrumental values are likely to be pursued by a broad spectrum of situated intelligent agents. (132)

 $^{^{42}}$ Of course, value drift or the possibility of such an artificial system being able to change its values implies that solving the value-alignment problem alone will not suffice to ensure safe and beneficial AI.

Thus, control problems are a type of value-alignment problem. The logical relation between these is that value alignment is necessarily prior to control: value alignment is a problem even for 'narrow' or 'weak' AI systems, whereas control is a problem only for more robust AI systems—e.g., artificial general intelligence or 'strong' AI.

To summarise, principal-agent problems arise whenever (1) an entity acts on behalf of another and (2a) their values are misaligned or (2b) there are informational asymmetries between them. In an economic context, the actors are typically described as individual (human) agents, government bodies, corporations, etc. However, this is not a logical requirement. The agent may be artificial, in which case the problem is typically referred to as a 'value-alignment problem'. The valuealignment problem is the problem of ensuring that an AI system is properly aligned with human values. This problem is functionally equivalent to the (more general) principal-agent problem from economics, which can arise in any situation where a principal delegates authority to an agent who acts on the principal's behalf. This specification of a principal-agent problem is the situation we find ourselves in with the increasing autonomy of AI systems today. Importantly, the functional equivalence of principal-agent and value-alignment problems implies that informational asymmetries are the main driver of the latter since value misalignment alone is neither necessary nor sufficient for the former. Finally, value-alignment problems are ubiquitous to AI systems, regardless of how narrow they are. Therefore, problems of value alignment are necessarily prior to problems of control, which arises only in the context of a sufficiently strong AI system.

2.3. Necessary Conditions. Before defending the main claim in Section 3, the final thing to clarify is the logical connection between linguistic communication and robust value alignment—i.e., *necessity*. Namely, the main claim says that linguistic communication (as described in Section 2.1) is *necessary* for robust value alignment between actors (as described in Section 2.2). 'Necessary' is meant in the bread-and-butter sense of the material conditional.⁴³ Contrapositively, one cannot achieve the robust form of value alignment required for safe and beneficial artificial intelligence without linguistic communication.

Thus, three logically equivalent ways of specifying the **main claim** quasi-formally in English are as follows:

(1) Main Claim.

Linguistic communication is *necessary* for (the possibility of) robust value alignment.

⁴³That is, $P \to Q$.

(2) Alternative Claim.

Robust value alignment (between actors) is possible *only if* those actors can communicate linguistically.

(3) Contrapositive Claim.

Without natural language, robust value alignment is not possible.

The equivalent formulation of the **main claim** that (2) proffers helps to clarify what is meant by 'robust'. Namely, we can understand value alignment as a matter of *degrees*—from perfectly misaligned to perfectly aligned. Thus, (2) underscores that *some* communicative ability (in the sense of information transfer discussed in Section 2.1) will be necessary for *some degree* of value alignment. This is discussed further in Section 3.1.

Therefore, 'robust' value alignment can be understood to specify a situation in which two actors' values (objectives) are sufficiently close to perfectly aligned. How close 'sufficiently close' needs to be for 'robust' value alignment will depend upon the context to which the values are relevant. Effectively, when the stakes are arbitrarily high, values will need to be arbitrarily close to perfectly aligned to say that the alignment of values is robust. In situations where the stakes are relatively low, it may be tolerable for values to be somewhat misaligned (again, relative to the context in question), but this would still be sufficient to claim robustness. Thus, the 'robust' component of robust value alignment is intentionally vague since the requirement will be highly sensitive to context.

However, part of the impetus for discussing value-alignment problems in the context of artificial intelligence is that the stakes appear to increase as these systems become more powerful—hence the importance of the **main claim**, defended in the subsequent section.

3. The Main Claim, Defended

Given the clarifications from Section 2 about what is meant by 'linguistic communication', 'robust value alignment', and the necessary connection between the two, we can now see why the **main claim** might be plausible. I provide two arguments in favour of the **main claim**. The first argument hinges on the unique advantages of linguistic communication for information transfer (Section 3.1). The second argument proceeds from the inherent rigidity of present-day objective functions and their formulation (Section 3.2). Neither of these arguments is definitive; however, I take them to underscore the plausibility of the **main claim**. That said, in Section 3.3, I provide additional empirical evidence in favour of the **main claim**. Section 4 examines some consequences if the **main claim** is true. 3.1. Language, Value Alignment, and Information Transfer. The first argument in favour of the main claim follows from the joint realisation that (1) value-alignment problems are structurally equivalent to principal-agent problems, discussed in Section 2.2.2; (2) principal-agent problems are fundamentally problems of information transfer, discussed in Section 2.2.1; and (3) linguistic communication, owing to features of systematicity and generalisability, is a uniquely robust and flexible communication system which allows for the transfer of information to an arbitrary degree of specificity, as discussed in Section 2.1.

Thus, the simplest way of understanding why we should expect the **main claim** to hold is as follows:

- (1) Principal-agent problems are primarily problems of informational asymmetries.
- (2) Value-alignment problems are structurally equivalent to principal-agent problems.
- (3) **Therefore**, value-alignment problems are primarily problems of informational asymmetries.
- (4) Any problem that is primarily a problem of informational asymmetries requires information-transferring capacities to be solved, **and** the more complex (robust) the informational burden, the more complex (robust) the information-transferring capacity is required.
- (5) **Therefore**, value-alignment problems require information-transferring capacities to be solved, **and** sufficiently complex (robust) value-alignment problems require robust information-transferring capacities to be solved.
- (6) Linguistic communication is a *uniquely* robust information-transferring capacity.
- (7) **Therefore**, linguistic communication is necessary for sufficiently complex (robust) value alignment.

Premise (1) was argued for in Section 2.2.1. Essentially, there are four possible combinations of values and information between a single principal and a single agent: either their values (objectives) are aligned, or they are misaligned, and either there are informational asymmetries between the actors, or there are not. So, we can ask which of the four combinations can possibly give rise to an agency dilemma.

Suppose the objectives of the principal and the agent are perfectly aligned, and there is no hidden information. In that case, no principal-agent problem is generated—the agent has the same values and information as the principal, meaning that if the agent acts in her own best interest, she will instantiate the principal's objectives by default. If objectives are misaligned, and there is perfect information, then there is *still* no problem generated, as we have seen.⁴⁴ It is only *possible* to generate an agency dilemma with imperfect information. As argued in Section 2.2.1, this can happen even when objectives are perfectly aligned. So, the information

⁴⁴Recall that if objectives are misaligned, but there is no private information, then 'the principal could propose a contract which perfectly controls the agent and induces the latter's actions to be what he would like to do himself in a world without delegation' (Laffont and Martimort, 2002).

component of the interaction generates the principal-agent problem, not the alignment or misalignment of objectives. The four possible situations are summarised in Table 1.

Va	lues
----	------

		Aligned	Misaligned
Information	Perfect	No	No
	Imperfect	Yes	Yes

TABLE 1. Combinations of information and values, giving rise to possible agency dilemmas (or not)

The key thing to note is that *informational asymmetries* give rise to principalagent problems, at least in an economic model. One may worry that this view conflates all incompleteness problems to information. The problem of misaligned objectives is dissolved when it is possible to write a complete contract, and informational asymmetries are fundamental in the real-world inability to specify complete contracts. However, it is also essential that arbitrary contracts include adequate incentives or penalties to properly shift the agent's objectives to align with the principal's objectives.⁴⁵ The question of shifting objectives is raised in Section 3.2; however, the important thing to note is that *understanding* the incentives or punishments is a prerequisite for incentives or punishments to shift objectives. Whether this understanding requires robust information transfer, in general, is questionable—rewards and punishments in basic reinforcement learning may align *simple* objectives via, e.g., food or a shock. However, as the environment and the values that agents carry in those environments become more complex, the difficulty of *specifying*—let alone aligning—these values increases exponentially. This can be seen in many real-world examples from AI research. Although reinforcement learning can be very useful for complex environments where objectives are difficult to define, it is also difficult to specify rewards; Sumers et al. (2022) highlight that linguistic communication 'is an intuitive and expressive way to communicate reward information to autonomous agents' (1). And, as was discussed in Section 2.2.2, step-wise progress toward an ultimate goal may lead to unexpected or undesirable behaviour.

Premise (2) was argued for in Section 2.2.2. As we have seen, the value-alignment problem for artificial intelligence is structurally equivalent to the principal-agent problem insofar as the *form* of the value-alignment problem arises from the dynamics of multi-agent interactions involving the delegation of tasks from one actor to another. 'The value-alignment problem', in the context of AI systems, is just a

⁴⁵Thanks to Gillian K. Hadfield for raising this point to me.

type of principal-agent problem where the agent is an artificial system. The subconclusion given in Premise (3) follows logically from Premises (1) and (2).

The first conjunct of Premise (4) should be relatively uncontroversial. If a problem fundamentally involves informational asymmetries, it requires *some* ability to transfer information for its solution. The second conjunct of Premise (4) suggests that the necessary resources for information transfer effectively scale with the complexity of the problem. As the problem of aligning values becomes more complex, coordination on those values requires more sophisticated information-transferring capacities on the part of the actors involved in the coordination problem.⁴⁶

In the language-origins literature, the *co*-evolution of coordination (cooperation) and communication explains why language *only* evolved in the hominin lineage (Sterelny, 2012; Planer and Sterelny, 2021). Considering the informational burden of robust cooperative capacities, Planer and Sterelny (2021) highlight that

the shift [in the hominin lineage] to delayed-return cooperation introduced a range of new social challenges. It is at this stage that tracking the reputation of third parties via gossip or its equivalent became essential to the stability of cooperation, which drove the need for more complex communicative technologies. (xiv)⁴⁷

So, I also take the second conjunct of Premise (4) to be relatively uncontroversial, although I will expand upon this in more detail in section 3.3. In any case, Premise (5) follows logically from (3) and (4).

Perhaps the most contentious of these Premises is (6). As was suggested in Section 2.1, the primary purpose of linguistic communication is communication i.e., information transfer. It is not the case that every feature of natural language is a requirement here. Nonetheless, the features of natural language that allow for robust information transfer are precisely the features that are understood to differentiate linguistic communication from simpler signalling systems—i.e., the compositional features of linguistic communication, among other things. These features give rise to systematicity and generalisability and, therefore, the flexibility and robustness of natural language for information transfer.

There are actually two claims here. First, communication systems are best understood in terms of information transfer, and second, linguistic communication systems are uniquely robust. Even though the referent of 'linguistic communication' is grossly underspecified, it is typically assumed that these capacities are unique to humans—namely, linguistic communication involves discrete, arbitrary,

 $^{^{46}}$ Using the signalling-game framework, Steinert-Threlkeld (2016) argues that compositional signalling is only evolutionarily beneficial when the world is sufficiently complex. It should be unsurprising that the structural properties of language are affected in non-trivial ways by the world in which those language evolve (Barrett and LaCroix, 2022).

 $^{^{47}}$ See also discussion in Sterelny (2014, 2021).

semantically meaningful signals whose referents may be displaced from the immediate environment (LaCroix, 2021). In linguistic communication, these signals are *structured* and *flexible*. Although *proto-languages* can involve displaced reference and other semantic features of words, the word-like sequences thus produced are not syntactically organised (Planer and Sterelny, 2021). Thus, researchers in linguistics, sociobiology, ethology, behavioural ecology, and related fields often take these claims for granted.⁴⁸

The conclusion, (7), is the **main claim**, which follows logically from Premises (5) and (6). Hence, understanding value-alignment problems and language in the ways I have specified in Section 2 logically entails that linguistic communication is necessary for robust value alignment.⁴⁹

3.2. **Objective Functions and Value Proxies.** Another way to think about why language might plausibly be required for a robust form of value alignment is to consider how the *symbolic systems* approach to AI—sometimes called 'good old-fashioned AI' (GOFAI)—is thought to have failed—or, perhaps more charitably, to be much more limited than initially believed.⁵⁰ Part of this is that these systems are far too rigid: every rule for action must be hard-coded. This is adequate for simple tasks, but as the complexity increases, it becomes intractable to write explicit instructions for every contingency.

Part of why 'second-wave' AI—ML and particularly DL methods bolstered by big data—has been surprisingly successful in comparison is that many rules for action are not hard-coded but implicit. What underwrites this intuition is that it is difficult to express the intuitive knowledge required for robust generalisation in

⁴⁸See, e.g., Otte (1974); Green and Marler (1979); Seyfarth et al. (1980a); Zahavi (1987); Hauser (1996); Smith (1997); Maynard Smith and Harper (2003); Fitch (2008); Bradbury and Vehrencamp (2011) and discussion in LaCroix (2020).

 $^{^{49}}$ One may object that it is possible that some future technologies may allow for direct interface with human thought, thus implying linguistic communication is not uniquely robust—direct information transfer from the source would also be sufficiently robust. In this case, it would be possible to 'read off' the principal's values directly, making it possible to align values without the need for linguistic communication. However, the Language of Thought Hypothesis (LOTH) proposes that thinking occurs in a 'mental language', sometimes called *mentalese* (Fodor, 1975). Assuming this is true, direct access to thought implies transferring information via a linguistic communication system, with the only distinction being modality. In the same way that it does not matter whether the transfer of information happens via written, spoken, or signed linguistic systems; it also does not matter whether the mode of linguistic communication is mental. Importantly, mentalese, like spoken, signed, or written natural language, is supposed to be compositional, implying that it has precisely the features I suggest are required for robust value alignment. The Principle of Compositionality (of Mental Representations) says that complex representations are composed of simple constituents, and the meaning of a complex representation depends upon the meanings of its constituents together with the constituency structure into which those constituents are arranged (Rescorla, 2019). See also Fodor (1987, 2008); Normore (1990).

 $^{^{50}}$ GOFAI is also called 'first-wave AI' or 'symbolic AI'. The phrase used here was coined by Haugeland (1985). For discussions of the failures of GOFAI, see also McDermott (1987); Cantwell Smith (2019).

the form of a set of verbally expressible rules that can be codified in a machine language. The symbolic systems approach of first-wave AI relied on the ability of humans to express explicit knowledge (often in the form of complicated *if-then* rules). In contrast, LaCroix and Bengio (2019) highlight that deep neural networks can 'capture' precisely the kind of implicit knowledge that is difficult to express in a formal language.⁵¹

The failures of symbolic systems are supposed to have been caused by the system's being 'brittle, unconducive to learning, defeated by uncertainty, and unable to cope with the world's rough and tumble' (Cantwell Smith, 2019). However, in contemporary approaches to AI, the 'values' (objective functions) encoded in the systems are what drive learning. These objective functions are still hard-coded. Thus, although these systems are significantly more flexible for *learning* and *acting*, the values encoded in these systems are still brittle. By analogy, it appears that a similar move from rigidity to flexibility concerning these systems are aligned with our values. In Section 3.3, I discuss, by analogy with humans, some empirical reasons why linguistic communication is necessary for robustly flexible values.

One may object that inverse reinforcement learning (IRL) furnishes a counterexample to this claim. It does not. Whereas a classic RL model is given a reward function and attempts to learn behaviour based on the rewards it receives, an IRL model is given behaviour and attempts to learn the reward function that would give rise to that behaviour.⁵² Therefore, the objective function is learned instead of hard-coded. So, the thought goes, it may be possible for a sophisticated IRL model to learn values from behaviours insofar as actions transfer information. However, this line of reasoning is problematic in several directions.

On the one hand, IRL is underspecified as a research problem because many reward functions can explain an actor's behaviour, and the costs of solving the problem tend to grow disproportionately with the size of the problem (Arora and Doshi, 2021). Essentially, many different reward functions could explain any observed behaviour.⁵³

On the other hand, behaviours are only ever going to serve as a proxy for values, which is part of why poorly specified objective functions give rise to value

⁵¹See also discussion in Buckner (2019).

⁵²Learning in this context may be through passive observation of behaviour (Ng and Russell, 2000; Abbeel and Ng, 2004), or it may involve active preference learning (Markant and Gureckis, 2014; Christiano et al., 2017; Basu et al., 2018).

 $^{^{53}}$ This is effectively identical to the problem of rule-following, discussed in Wittgenstein (1953); Kripke (1982) and a significant body of secondary literature since.

misalignment in the first place, as discussed in Section 2.2.2. Furthermore, by emphasising the information transferred by behaviours alone, we run into the well-known problems to which revealed preference theory from economics gives rise.⁵⁴ An IRL model's basic (and false) assumption is that the behaviour it observes *is* optimal, given the (hidden) reward function motivating that behaviour. In addition, even if behaviours were a good proxy for preferences and provided a high-fidelity channel for transferring information about those preferences, linguistic communication would aid learning objectives at a much higher rate. The rate will be highly relevant for certain aspects of *robust* value alignment. In addition to systematicity and generalisability, linguistic communication offers the possibility of speed to value alignment. In high-stakes cases, it will be impossible to wait until a sufficient number of behavioural instances are observed.

3.3. The Curious Case of Human Value Alignment. As we have seen, the value-alignment problem for artificial intelligence asks how we design the *objective functions* of AI systems to guarantee maximal overlap between the actions of the artificial agent and the objectives or values of (human) principals. Despite the structural identity of value-alignment problems for AI systems and the principal-agent problem more generally, as I mentioned in Section 2.2.2, the solutions to these problems may be quite distinct. This point is made clearer when we realise that principal-agent problems between humans allow for assumptions that are not warranted when the agent is an artificial system—namely, the ability to communicate one's *intended meaning* to an arbitrary degree of specificity via natural language.

When considering problems in artificial intelligence, it is often instructive to ask, *How do we do it?* For the entire history of the species, humans have engaged in the production of agents with human-level intelligence, whose values may be misaligned with their own, and over which they have limited control—i.e., human children. This provides something of a *proof of concept* that values can be aligned in such agents (Christian, 2020). From the cooperative nature of our species, it seems apparent that humans are at least capable of aligning their values in the variety of contexts we face daily. Part of the reason we can do so is that we can communicate via natural language. Humans use linguistic competence to impart extremely subtle norms, goals, and values in subsequent generations that align with a long cultural history of norms, goals, and values. Because human linguistic capacity comes hardwired, we can take this particular aspect of how we align our values for granted. Thus, the blind spot of value-aligned agency arises from a lack of sensitivity to the importance

⁵⁴Revealed preference theory (Samuelson, 1938a,b) uses consumer behaviour to analyse the choices made by individuals. See discussion in Sen (1973, 1977, 1993, 1997, 2002); Koszegi and Rabin (2007); Hausman (2012).

of language for *Homo sapiens* as fundamentally cooperative and social creatures.⁵⁵ What is pressing in the context of value alignment for AI is that assumptions of linguistic competency cannot be taken for granted when considering artificial agents.

Part of why language evolved in *Homo sapiens* is because of the cooperative demands that evolved in our lineage. In the context of hominin evolution, cooperation would have included 'quite demanding forms of collective action; and with greater cooperation comes greater communication' (Planer and Sterelny, 2021, 73). It is difficult to imagine how such robust cooperation could have evolved without corresponding increases in the complexity and flexibility of our communicative abilities. For example, the control of fire (it is supposed) would require communication systems much more complex than anything seen in other great ape species (though not necessarily language) because controlling fire poses a range of sophisticated coordination and cooperation problems.⁵⁶

One can understand the intent of other humans precisely because one can communicate linguistically. Strong empirical evidence exists for a tight, bidirectional connection between theory of mind capacities and linguistic capacities in human infants (Astington and Jenkins, 1999; Astington and Baird, 2005; de Villiers, 2007; de Villiers and de Villiers, 2014). First, language appears necessary to learn to express concepts surrounding one's own feelings and inner world (Nelson, 2005; Dunn and Brophy, 2005; Hutto, 2012); second, the information that language conveys about others' thoughts, feelings, beliefs, desires, etc., is much richer than the information conveyed through behaviour, eye gaze, gestural expressions, etc. (Appleton and Reddy, 1996; Harris, 2005; Peterson and Siegal, 1999; Wellman and Peterson, 2013); third, linguistic abilities allow individuals to reason abstractly about others' actions via their beliefs (Astington and Jenkins, 1999; de Villiers and de Villiers, 2009; Milligan et al., 2007).

⁵⁵Of course, one might object that although humans use linguistic communication systems—i.e., they have linguistic communicative abilities—they often fail to align their values. This is fine. The **main claim** furnished a *necessary* condition for value alignment; it does not say anything about whether language is *sufficient* for value alignment. It should be obvious that it is not. There is a substantive question about whether any set of abilities will suffice for value alignment; however, my claim is that no matter what set of abilities is uncovered, value alignment will be impossible without linguistic communication. The **main claim** also does not imply linguistic communication is *the only* necessary condition for value alignment.

 $^{^{56}}$ See, e.g., Ofek (2001); Gowlett (2006, 2016); Garde (2009); Gowlett and Wrangham (2013); Pickering (2013); Gamble (2013) and discussion in Twomey (2013, 2014); Planer and Sterelny (2021).

Further, the compositional nature of language allows humans to communicate their goals with one another to an arbitrary degree of specificity.⁵⁷ Linguistic communicative abilities also affect the cultural accumulation of new concepts. High-fidelity cultural learning allows human populations to solve coordination/cooperation problems because it allows selective learning and the accumulation of small improvements over time (Boyd et al., 2011). That is, language allows for accumulating knowledge across generations via social learning.⁵⁸ Planer and Sterelny (2021) high-light that 'some forms of cooperation are stable only if reputation (knowledge of the past social actions of others) is tracked reliably and is part of common knowledge' (25). Again, this robust sort of cooperation depends significantly upon language.

Communication depends, to some extent, upon joint attention and common knowledge (Tomasello, 2008, 2014). Complex language is not necessary for joint commitment if there is common-ground understanding (Tomasello, 2014); however, diminishing common ground between agents appears to necessitate greater lexical and structural richness in the language used to communicate information about disparate knowledge between agents.⁵⁹ Essentially, when the social aspects of agent interactions become increasingly dispersed in time and space, members of a social group will need more sophisticated communication (and cognitive) abilities to report behaviour and events that happened 'elsewhere and elsewhen' (Planer and Sterelny, 2021, 195).

Furthermore, human *moral* behaviour is rooted in uniquely-human linguistic abilities precisely because of the cultural accumulation that language 'exemplifies and enables' (Bryson, 2018, 19).⁶⁰ Poulshock (2006) suggests that once complex language exists, humans have access to a new type of cultural evolution that can aid in promoting altruistic behaviour. A linguistically-encoded, abstract value system can promote and regulate moral behaviour to the extent that a sufficient number of members participating in such a moral system provides indirect benefits through, e.g., protection and group support. Language, therefore, allows humans to overcome genetic self-interest in a way not seen in other species in nature, despite a lack of direct benefits for moral behaviour. This relates to the importance of linguistic communication to delayed-return cooperation mentioned above (Section 3.1). Language allows social groups to encode moral principles and to demand adherence

26

⁵⁷In addition, in the context fo AI, it has been suggested that linguistic communication may be necessary for transparency, trust, and explainability. For example, Williams (2018) argues that an ethical AI system that is asked for clarification may unintentionally communicate that it would be willing to perform an unethical action (even if it is unable to do so). Further, this type of miscommunication may negatively influence the morality of its human teammates. This point is discussed further in Jackson (2019); Jackson and Williams (2019).

⁵⁸See, e.g., Boyd (2016); Henrich (2016); Planer and Sterelny (2021).

 $^{^{59}}$ This been observed empirically in, e.g., children's sign languages; see Meir et al. (2010).

⁶⁰See also discussion in Bryson (2008); Malle and Scheutz (2014); Malle (2015, 2016).

to those principles. It can turn high-cost behaviour into low-cost behaviour, thus aiding the enforcement of moral principles or norms (Sober and Wilson, 1998). Furthermore, such moral codes can be transmitted across generations via spoken or written language.

Language also provides a cost-efficient means for classifying certain behaviours as good or bad. It allows individuals to tag others as immoral (defectors/cheaters) or as moral (cooperators) via gossip, blame, praise, etc., when they adhere or fail to adhere to these principles; namely, linguistic communication allows us to articulate norms and monitor behaviour (Boehm, 2000). Thus, linguistic ability significantly affects *moral* ability because it allows for a simple, efficient, systematic, generalisable, flexible, and (typically) high-fidelity means of transmitting information.

Empirical studies appear to corroborate the claim that moral behaviour depends (at least to some extent) on linguistic ability. For example, de Waal (1996) suggests that the following (pre-)conditions are necessary for the evolution of morality: group value, mutual aid, and internal conflict. Actors must then resolve intragroup conflicts to balance individual and collective interests, either at the dyadic level (oneon-one interactions) or at a higher level. The latter may depend on community concern, which can be expressed in terms of mediated reconciliation or peaceful arbitration of disputes—as observed in many primate species, including chimpanzees. However, de Waal (1996) notes that higher-level interactions may also be expressed in terms of indirect reciprocity (appreciation of altruistic behaviour at a group level) or encouragement of contributions to the quality of the social environment; importantly, these last two may be limited to *human* moral systems (34). But, it should be relatively obvious how language can support these functions—via, e.g., gossip, blame, praise, or encouragement.

Considering the development of moral behaviour in humans, Several studies have suggested that moral judgements in children develop alongside language (Smetana and Braeges, 1990). Infants appear to be prosocial, actively helping others from around 14 months of age (Warneken and Tomasello, 2006, 2007). Zahn-Waxler et al. (1992) report that prosocial behaviours, such as sharing, providing comfort, etc., emerge in human children between years 1 and 2. These behaviours increase in frequency and variety over this period. However, Although human children may act in helpful ways and display empathic reactions by 2 years, categorical judgements do not appear to emerge until around 3 years of age.⁶¹ This is also around the time that, a 'normative sense of obligation' begins to emerge (Tomasello, 2018, 250). So, the distinction between moral and mere social norms, or conventions, appears

⁶¹See discussion in Dahl (2018), and studies in Zahn-Waxler et al. (1992); Eisenberg et al. (2007); Roth-Hanania et al. (1976); Marco F. H. Schmidt and Hannes Rakoczy and Michael Tomasello (2012); Smetana et al. (2012); Dahl and Kim (2014); Dahl (2015).

to arise around three years.⁶² Importantly, for our purposes, these behaviours appear to develop and become more complex in parallel with linguistic development timelines.⁶³ Non-human primates, who lack linguistic capacities, display only *proto*-moral behaviours.⁶⁴

All these considerations are couched in the language of *morality*, which, like linguistic communication, is often considered uniquely human. One difficulty with such discussions is that it is not always clear what is meant by 'morality' or what constitutes moral behaviour (Dahl, 2018).⁶⁵ However, this is stronger than what is required for the **main claim**. Aligning values is a type of coordination problem. Korsgaard (1996) suggests that the publicity of language allows values to be shared.⁶⁶ So, at a structural level, we can understand value alignment in terms of coordination and cooperation in addition to the problem of incentive structures. It is well understood that certain classes of coordination problems benefit from cheap talk (Farrell and Rabin, 1996)—i.e., simple communication. Therefore, simple communication channels will allow for simple alignment of values; however, robust (linguistic) communication is necessary to ensure robust value alignment in complex environments.

To summarise, the first argument in favour of the **main claim** depends upon the conceptualisation of language and value alignment, discussed in Section 2. The **main claim** follows logically from how these concepts are specified. The key constitutive features of linguistic communication—simplified here in terms of

⁶²This distinction is more firmly established by four years, paralleling linguistic development in children. Early models of moral development in human children (Piaget, 1923, 1924, 1932; Kohlberg, 1958, 1963, 1964, 1969, 1981, 1984) suggest a hierarchical development through several stages of understanding moral concepts—pre-conventional, conventional, and post-conventional from pre-adolescence to adulthood. However, more recent studies have suggested that moral development occurs much earlier. It is also worth noting that the validity of these classical models has been called into question for several reasons, including possible gender bias (Gilligan, 1982). ⁶³See extended discussion in Lust (2006). Of course, progression rates will vary highly between individual children (Brown, 1973).

 $^{^{64}}$ There is little evidence of instrumental helping in chimpanzees, although there is a *tendency* for helping under certain conditions—for example, chimpanzees are more likely to help when the task does not involve food, when information about the goal can be transferred through a non-linguistic communication channel, and when the recipient of help is not a conspecific (Warneken and Tomasello, 2006).

 $^{^{65}\!\}mathrm{See}$ also discussion in Greene (2007); Wynn and Bloom (2014).

⁶⁶In her view, *reasons* are essential to morality (which she approaches from a distinctly Kantian point of view). Part of this involves moral reflection from the first-personal perspective. In analogy with considerations concerning the impossibility of a private language (Wittgenstein, 1953), Korsgaard thinks that there are no *private reasons*—reasons have a public normative force similar to how language has a public normative meaning. This view is somewhat complicated because it relies on Korsgaard's constructivist views surrounding identity; in any case, see discussion in (Pauer-Studer, 2003), but cf. Gibbard (1999).

compositionality—give rise to the systematicity and generalisability that is a prerequisite for robust value alignment, understood as fundamentally a problem of information transfer. The second argument for the **main claim** highlights that objective functions are rigid in precisely the ways that contributed to the failures of the symbolic systems approach to AI. Hence, although current systems are flexible regarding learning, they are still rigid regarding value alignment. Linguistic communication is necessary for flexible objective specification, which is necessary for robust value alignment. Finally, the connection between language and value alignment was explored in the context of cooperation in the hominin lineage. This provides additional empirical evidence of the necessity of language for robust value alignment while simultaneously highlighting the fact that we cannot take these capacities for granted when considering how to ensure that the values of AI systems are aligned with our own values.

In the next section, I explore some consequences following from the main claim.

4. DISCUSSION

What follows if the **main claim** is true? At minimum, the **main claim** specifies a demanding lower bound on the difficulty of solving the value-alignment problem.

At the same time, Bostrom (2014) suggests that achieving fully human-level performance on natural language is an 'AI-complete' problem. This means that creating linguistic AI is essentially equivalent in difficulty or complexity to creating generally human-level AI systems (17). This implies that if one were to create a system fully capable of linguistic communication, in all likelihood, one would 'also either already have succeeded in creating an AI that could do everything else that human intelligence can do, or they would be but a very short step from such a general capability' (17).⁶⁷

If Bostrom (2014) is correct about language being an AI-complete problem, then linguistic AI is functionally equivalent to human-level AI. If the **main claim** is true, then linguistic AI is, in some sense, prior to robust value alignment. But, it was suggested above (2.2.3) that value alignment is logically prior to problems of control.⁶⁸ The conjunction of these three insights leads to circularity, which may imply that solving robust value alignment problems cannot be achieved until *after* we have already created an AI system that we are effectively unable to control; see Figure 1.

 $^{^{67}}$ Though, cf. Goebel (2008).

⁶⁸Recall that even narrow AI systems with limited functionality can fail to do what we intended, in the domain to which it is limited. In contrast, control problems arise in the context of stronger AI systems, on a par with human-level intelligence or superintelligence.



FIGURE 1. Circularity following from the **main claim** in conjunction with AI Completeness

This implication depends on a nested conditional: *If* language is an AI-complete problem, *then if* the **main claim** is true, robust value alignment may be impossible at least before the creation of an uncontrollable AI system. If any of the antecedents is false, then this may open the door to some (cautious) optimism about our ability to solve the value-alignment problem prior to the advent of a system that we cannot control. Fortunately, AI completeness is an *informal* concept, defined by *analogy* with complexity theory; so, it is possible that Bostrom (2014) is incorrect.⁶⁹ This leaves room for work on alignment with linguistic, albeit controllable, AI systems.

In addition, the principal-agent model is highly idealised, as all models are. Hence, it may be that the **main claim** depends upon an idealised (hence false) assumption. Still, it is useful for underscoring that informational asymmetries, rather than value misalignment *per se*, are the driving force of value-alignment problems. This insight opens the door to my argument that linguistic communication is at least a highly effective—if not necessary—means for transferring information to an arbitrary degree of specificity, which would allow on-the-spot adjustments in action in a way that is difficult to imagine accomplishing without language.⁷⁰

Counter-intuitively, the main driver of the value-alignment problem is *not* misaligned values but informational asymmetries. Thus, value-alignment problems are primarily problems of aligning *information*. However, this is not to say that values play no role in solving value-alignment problems. Information alone cannot *cause*

⁶⁹Although, see discussion in Shahaf and Amir (2007).

⁷⁰Consider that any solution to a significantly complex value-alignment problem, driven by informational asymmetries, will require some degree of generalisability to accommodate unseen instances; and, that solution will require some degree of systematicity to accommodate efficiency constraints. As a result, such a solution will be inherently linguistic.

action. Even in the classic economic and game-theoretic contexts, values are not objective. As Hausman (2012) highlights, a game *form*, complete with (objective) payoffs which correspond to states of the world conditional on actions, does not specify a *game* because it does not specify how the agents *value* the outcomes of the game form. For a game to be defined, it is not sufficient to know the objective states that would obtain under certain actions; instead, one must also know the *preferences* (values) of the actors over those (objective) outcomes. Thus, a specification of results needs to coincide with 'how players understand the alternatives they face and their results' (Hausman, 2012, 51).

What comes out of the discussion of value alignment qua principal-agent problems is that the main issue involves aligning values by ensuring appropriate *incen*tive structures. This is true in economics, because both the agent and principal are assumed to be utility maximisers. The principal must limit the divergence of the possible actions of the agent from the principal's interests—for example, by introducing incentives for the agent or paying some costs for monitoring to limit the agent's (possibly unaligned) actions. However, although principal-agent and valuealignment problems arise for the same reasons, there are dissimilarities in how they can be solved. For example, we can take certain things for granted when we consider interactions between human principals and human agents. In addition to linguistic communication abilities, human agents come with in-built values, whereas AI systems have their 'values' programmed in an *objective function*. Given the complexity of accurately describing our objectives in a programming language, it is vanishingly unlikely that we can succeed in specifying the 'correct' objective function for an artificial system—if such a specifiable function even exists—as was discussed in Section 3.2.

One may object that 'large language models' (LLMs)—e.g., BERT (Devlin et al., 2019), GPT-3 (Brown et al., 2020), RETRO (Borgeaud et al., 2022), LaMDA (Thoppilan et al., 2022), etc.—furnish a counterexample to the **main claim** (and therefore its consequences). The thought is that LLMs *are* linguistic. However, this is misguided. The phrase 'large language model' is a misnomer better described in terms of 'large corpus models' (Veres, 2022) or 'stochastic parrots' (Bender et al., 2021). Some authors argue that meaning cannot be learned from form alone (Bender and Koller, 2020; Bisk et al., 2020; Marcus and Davis, 2020); in this sense, LLMs do not capture meaning and, therefore, cannot be linguistic.⁷¹

Although LLMs can be impressive syntax engines, language is more than mere syntax (Section 2.1). Importantly, language is a primarily *social* endeavour, which requires the participation of both senders and receivers. Part of the reason that

 $^{^{71}}$ See, e.g., McCoy et al. (2019); Ettinger (2020); Pandia et al. (2021); Sinha et al. (2021); Sahlgren and Carlsson (2021).

the output of LLMs *looks* so impressive is that it is being interpreted by competent language users (humans) who project meaning, intentions, understanding, etc., onto the system.

To illustrate this, prompting DALL-E—a version of GPT-3 that generates images from text inputs—to produce a 'self-portrait' results in several images of apparently white, male humans; See figure 2. It would be strange to suggest from this that



FIGURE 2. Output of DALL-E mini given prompt 'DALL-E Self Portrait'. See https://huggingface.co/spaces/dalle-mini/dalle-mini

DALL-E 'thinks' that *it* is a white, male human—a projection from the user's *interpretation* that a self-portrait applies (in fact) to *oneself*. What actually happens is this. DALL-E is trained on a dataset of text-image pairs. The output merely indicates that 'self-portrait' is typically paired with images of white, male humans. There is nothing more to this than mere correlation, which is, at bottom, all that underlies present-day AI systems.

5. Conclusion

The value-alignment problem is ubiquitous in the context of AI systems. As these systems become more sophisticated, these problems become more pressing. The **main claim** helps clarify how difficult value alignment can be in this context. However, empirical evidence proves that aligning values between agents with human-level intelligence is at least *possible*. That said, it is important to maintain sensitivity to the informational asymmetries that underlie value alignment for AI systems since value alignment between human principals and human agents allow key features of the problem to be taken for granted—particularly, the ability to communicate linguistically. No such assumption can be made when the agent is artificial.

It is precisely because we cannot make assumptions about, e.g., in-built linguistic ability that the implementation of ethically-aligned artificial systems comes to bear on normative theory. Normative theories historically focus on *human* agents. As such, they take for granted that agents can communicate linguistically. The creation of artificial systems allows for no such assumptions. As a result, thinking about problems of value alignment in terms of coordination—and linguistic aids to successful coordination—provides valuable insights into the implicit foundations upon which many normative theories rest.

Thinking about the importance of linguistic communication for value alignment in artificial systems pushes work in this area beyond the straightforward application of specific insights from normative theory to AI by additionally providing novel insights in the opposite direction. Importantly, such insights will be theory-neutral insofar as they do not depend upon any particular metaethical framework.

Hence, linguistic communication has been a blind spot in theorising about value alignment in the context of artificial agents; however, since the ability to communicate linguistically is taken for granted in normative theories that focus on human (natural) agents, little attention has been paid to the necessity of language for value-aligned agency in natural agents as well.

As mentioned at the outset, I take the **main claim** to hold in the context of value alignment more generally than just situations involving artificial agents. Hence, the linguistic blind spot of value-aligned agency applies in normative context for both natural and artificial agents.

Acknowledgements. I would like to thank the organisers of the Normative Theory and AI Workshop (2022)—Seth Lazar, Pamela Robinson, Claire Benn, and Todd Kharu—as well as participants—David Grant, Jeff Behrends, John Basl, Chad Lee-Stronach, Jack Parker, and David Danks—for helpful feedback. I would also like to thank, in no particular order, Aydin Mohseni, Jennifer Nagel, Yoshua Bengio, Dominic Martin, Benjamin Wald, Aaron Courville, Sasha Luccioni, Gillian Hadfield, Celso Neto, Duncan McIntosh, Aaron Wright, Jeffrey Barrett, Simon Huttegger, Richmond Campbell, and Michael Noukhovitch, for discussion and feedback throughout the process. This paper was presented at the Philosophy Department Colloquium at Dalhousie University (2021). I am grateful to the audience members for helpful discussion. Thanks also to the Schwartz Reisman Institute at the University of Toronto for partially funding this research, and to Mila - Québec Artificial Intelligence Institute for providing generous resources.

References

- Abbeel, Pieter and Andrew Y. Ng (2004). Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the 21st International Conference on Machine Learning (ICML 2004)*, pages 1–8. Association for Computing Machinery, New York.
- Akerlof, George A. (1970). The market for 'lemons': Quality uncertainty and the market mechanism. Quarterly Journal of Economics, 84(3): 488–500.
- Allen, Colin, Wendell Wallach, and Iva Smit (2006). Why machine ethics? IEEE Computer Society, 21(4): 12–17.
- Amodei, Dario and Jack Clark (2016). Faulty reward functions in the wild. https://openai.com/blog/faulty-reward-functions/.
- Amodei, Dario, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané (2020). Concrete problems in ai safety. arXiv preprint arXiv:1606.06565, pages 1–29. https://arxiv.org/abs/1606.06565.
- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner (2016). Machine bias. ProPublica, May 23: np. https://www.propublica.org/article/ machine-bias-risk-assessments-in-criminal-sentencing.
- Appleton, Michele and Vasudevi Reddy (1996). Teaching three-year-olds to pass false belief tests: A conversational approach. *Social Development*, 5(3): 275–291.
- Arbib, Michael A. (2002). The mirror system, imitation, and the evolution of language. In Nehaniv, C. and K. Dautenhahn, editors, *Imitation in Animals and Artifacts*, pages 229–280. The MIT Press, Cambridge, MA.
- Arbib, Michael A. (2003). The evolving mirror system: a neural basis for language readiness. In Christiansen, M. H. and S. Kirby, editors, *Language Evolution*, pages 182–200. Oxford University Press, Oxford.
- Arbib, Michael A. (2005). An action-oriented neurolinguistic framework for the evolution of protolanguage. In Tallerman, M., editor, *Language Origins: Perspectives* on Evolution, pages 21–47. Oxford University Press, Oxford.
- Aristotle (1995). History of animals, book IX. In Barnes, Jonathan, editor, The Complete Works of Aristotle, The Revised Oxford Translation, volume 1, pages 214–256. Princeton University Press, Princeton.
- Arnold, Kate and Klaus Zuberbühler (2006a). The alarm-calling system of adult male putty-nosed monkeys, cercopithecus nictitans martini. Animal Behaviour, 72: 643–653.
- Arnold, Kate and Klaus Zuberbühler (2006b). Language evolution: Semantic combinations in primate calls. *Nature*, 441: 303.
- Arnold, Kate and Klaus Zuberbühler (2008). Meaningful call combinations in a non-human primate. *Current Biology*, 18: R202–R203.
- Arnold, Kate and Klaus Zuberbühler (2013). Female putty-nosed monkeys use experimentally altered contextual information to disambiguate the cause of male alarm calls. *PLOS ONE*, 8(6): e65660.
- Arora, Saurabh and Prashant Doshi (2021). A survey of inverse reinforcement learning: Challenges, methods and progress. Artificial Intelligence, 297: 1–28.
- Arrow, Kenneth J. (1963). Uncertainty and the welfare economics of medical care. American Economic Review, 53: 941–973.
- Arrow, Kenneth J. (1968). The economics of moral hazard: Further comment. American Economic Review, 58: 537–539.

- Astington, Janet Wilde and Jodie A. Baird (2005). Introduction: Why language matters. In Astington, J. W. and J. A. Baird, editors, *Why language matters for* theory of mind, pages 2–25. Oxford University Press, Oxford.
- Astington, Janet Wilde and Jennifer M. Jenkins (1999). A longitudinal study of the relation between language and theory of mind development. *Developmental Psychology*, 35(5): 1311–1320.
- Barrett, Jeffrey A. and Travis LaCroix (2022). Epistemology and the structure of language. *Erkenntnis*, 87: 953–967. https://doi.org/10.1007/ s10670-020-00225-4.
- Basu, Chandrayee, Mukesh Singhal, and Anca D. Dragan (2018). Learning from richer human guidance: Augmenting comparison-based learning with feature queries. In Kanda, Takayuki, Selma Ŝabanović, Guy Hoffman, and Adriana Tapus, editors, 13th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pages 132–140, New York. Association for Computing Machinery.
- Belanger, Rachelle M. and Lynda D. Corkum (2009). Review of aquatic sex pheromones and chemical communication in anurans. *Journal of Herpetology*, 43: 184–191.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell (2021). On the dangers of stochastic parrots: Can language models be too big?
 In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pages 610–623.
- Bender, Emily M. and Alexander Koller (2020). Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. In Jurafsky, Dan, Joyce Chai, Natalie Schluter, and Joel Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Berwick, Robert C. and Noam Chomsky (2016). Why Only Us: Language and Evolution. The MIT Press, Cambridge, MA.
- Bickerton, Derek (1990). Language and Species. University of Chicago Press, Chicago.
- Bickerton, Derek (1995). Language and Human Behavior. University of Washington Press, Seattle.
- Bishop, John (1989). Natural Agency: An Essay on the Causal Theory of Action. Cambridge University Press, Cambridge.
- Bisk, Yonatan, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian (2020). Experience grounds language. In Webber, Bonnie, Trevor Cohn, Yulan He, and Yang Liu, editors, *Proceedings* of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 8718–8735, Online. Association for Computational Linguistics.
- Boehm, Christopher (2000). Conflict and the evolution of social control. *Journal* of Consciousness Studies, 7(1–2): 79–101.
- Borgeaud, Sebastian, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman

Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and Laurent Sifre (2022). Improving language models by retrieving from trillions of tokens. *arXiv preprint*, 2112.04426: 1–43. https://arxiv.org/abs/2112.04426.

- Bostrom, Nick (2014). Superintelligence: Paths, Dangers, Strategies. Oxford University Press, Oxford.
- Boyd, Robert (2016). A Different Kind of Animal: How Culture Made Humans Exceptionally Adaptable and Cooperative. Princeton University Press, Princeton.
- Boyd, Robert, Peter J. Richerson, and Joseph Henrich (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences of the United States of America*, 108(Suppl. 2): 10918–10925.
- Bradbury, Jack W. and Sandra L. Vehrencamp (2011). *Principles of Animal Communication*. Sinauer, Sunderland, MA, 2 edition.
- Brand, Myles (1984). Intending and Acting: Toward a Naturalized Action Theory. The MIT Press, Cambridge, MA.
- Bratman, Michael E. (1987). Intention, Plans, and Practical Reason. Harvard University Press, Cambridge, MA.
- Brown, Roger (1973). A First Language: The Early Stages. Harvard University Press, Cambridge, MA.
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei (2020). Language models are few-shot learners. arXiv preprint arXiv:2005.14165. https://arxiv.org/abs/2005.14165.
- Bruschini, Claudia, Rita Cervo, and Stefano Turillazzi (2010). Pheromones in Social Wasps. *Vitamins and Hormones*, 83: 521–549.
- Bryson, Joanna J. (2008). Embodiment versus memetics. *Mind & Society*, 7(1): 77–94.
- Bryson, Joanna J. (2018). Patiency is not a virtue: the design of intelligent systems and systems of ethics. *Ethics and Information Technology*, 20: 15–26.
- Buckner, Cameron (2019). Deep learning: A philosophical introduction. Philosophy Compass, 14: e12625.
- Cantwell Smith, Brian (2019). The Promise of Artificial Intelligence: Reckoning & Judgment. The MIT Press, Cambridge, MA.
- Cave, Stephen (2017). Intelligence: A history. *Aeon.* https://aeon.co/essays/ on-the-dark-history-of-intelligence-as-domination.
- Chomsky, Noam (1965). Aspects of the Theory of Syntax. The MIT Press, Cambridge, MA.
- Chomsky, Noam (1980). Rules and Representations. Basil Blackwell, London.
- Chomsky, Noam (1995). *The Minimalist Program*. The MIT Press, Cambridge, MA.
- Christian, Brian (2020). The Alignment Problem: Machine Learning and Human Values. W. W. Norton & Company, New York.

- Christiano, Paul F., Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei (2017). Deep reinforcement learning from human preferences. In Guyon, I., U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and
- R. Garnett, editors, Advances in Neural Information Processing Systems, pages 1–9, Red Hook, NY. Curran Associates, Inc.
- Costa-Leonardo, Ana Maria and Ives Haifig (2010). Pheromones and exocrine glands in Isoptera. *Vitamins and Hormones*, 83: 521–549.
- Creel, Kathleen (2021). Opening the Black Box: Explanation and Transparency in Machine Learning. PhD thesis, University of Pittsburgh.
- Creel, Kathleen A. (2020). Transparency in Complex Computational Systems. Philosophy of Science, 87(4): 568–589.
- Culicover, Peter W. and Ray Jackendoff (2005). *Simpler Syntax*. Oxford University Press, Oxford.
- Dahl, Audun (2015). The developing social context of infant helping in two us samples. *Child Development*, 86: 1080–1093.
- Dahl, Audun (2018). New beginnings: An interactionist and constructivist approach to early moral development. *Human Development*, 61(4–5): 232–247.
- Dahl, Audun and Lizbeth Kim (2014). Why is it bad to make a mess? preschoolers' conceptions of pragmatic norms. *Cognitive Development*, 32: 12–22.
- Davidson, Donald (1963). Actions, reasons, and causes. *The Journal of Philosophy*, LX(23): 685–700.
- Davidson, Donald (1971). Agency. In Marras, Ausonio, R. N. Bronaugh, and Robert W. Binkley, editors, Agent, Action, and Reason, pages 1–37. University of Toronto Press, Toronto.
- De Beule, Joachim (2008). Compositionality, Hierarchy and Recursion in Language: A Case Study in Fluid Construction Grammar. PhD thesis, Vrije Universiteit Brussel.
- de Boer, Jelle (2011). Moral ape philosophy. Biology & Philosophy, 26(6): 891–904.
- de Villiers, Jill G. (2007). The interface of language and theory of mind. *Lingua*, 117(11): 1858–1878.
- de Villiers, Jill G. and Peter A. de Villiers (2009). Complements enable representation of the contents of false beliefs: Evolution of a theory of theory of mind. In Foster-Cohen, S., editor, *Language acquisition*, pages 169–195. Palgrave Macmillan, Hampshire.
- de Villiers, Jill G. and Peter A. de Villiers (2014). The role of language in theory of mind development. *Topics in Language Disorders*, 34(4): 313–328.
- de Waal, Frans (1996). Good Natured: The Origins of Right and Wrong in Humans and Other Animals. Harvard University Press, Cambridge, MA.
- de Waal, Frans (2006). *Primates and Philosophers*. Princeton University Press, Princeton.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint, 1810.04805: 1–16. https://arxiv.org/abs/1810.04805.
- Dretske, Fred (1988). *Explaining Behavior: Reasons in a World of Causes*. The MIT Press, Cambridge, MA.
- Dunn, Judy and Marcia Brophy (2005). Communication, relationships and individual differences in children's understanding of mind. In Astington, J. W. and J. A. Baird, editors, Why language matters for theory of mind, pages 50–69. Oxford

University Press, Oxford.

- Ecoffet, Adrien, Jeff Clune, and Joel Lehman (2020). Open questions in creating safe open-ended ai: Tensions between control and creativity. *arXiv preprint*, 2006.07495: 1–9. https://arxiv.org/abs/2006.07495.
- Eisenberg, Nancy, Richard A. Fabes, and Tracy L. Spinrad (2007). Prosocial development. In Lamb, M. E. and R. M. Lerner, editors, Handbook of Child Psychology and Developmental Science: Volume III. Social, Emotional, and Personality Development, pages 610–656. Wiley & Sons, Hoboken, NJ.
- Eisenhardt, Kathleen M. (1989). Agency theory: An assessment and review. The Academy of Management Review, 14(1): 57–74.
- Enç, Berent (2003). How We Act: Causes, Reasons, and Intentions. Oxford University Press, Oxford.
- Erasmus, Adrian, Tyler D. P. Brunet, and Eyal Fisher (2021). What is interpretability? *Philosophy & Technology*, 34: 833–862.
- Ettinger, Allyson (2020). What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *arXiv preprint*, 1907.13528: 1–20. https://arxiv.org/abs/1907.13528.
- Falbo, Arianna and Travis LaCroix (2021). Est-ce que vous compute? codeswitching, cultural identity, and ai. arXiv preprint, 2112.08256: 1–19. https: //arxiv.org/abs/2112.08256.
- Farrell, Joseph and Matthew Rabin (1996). Cheap talk. The Journal of Economic Perspectives, 10(3): 103–118.
- Fitch, W. Tecumseh (2008). Kin Selection and 'Mother Tongues': A Neglected Component in Language Evolution. In Oller, D. K. and U. Griebel, editors, *Evolution of Communication Systems: A Comparative Approach*, pages 275–296. MIT Press, Cambridge, MA.
- Fitch, W. Tecumseh (2010). *The Evolution of Language*. Cambridge University Press, Cambridge.
- Fitch, W. Tecumseh (2017). Empirical approaches to the study of language evolution. *Psychonomic Bulletin & Review*, 24: 3–33.
- Fitch, W. Tecumseh, Marc Hauser, and Noam Chomsky (2005). The evolution of the language faculty: clarifications and implications. *Cognition*, 97(2): 179–210.
- Fitzpatrick, Simon (2017). Animal morality: What is the debate about? *Biology* and *Philosophy*, 32: 1151–1183.
- Flack, Jessica C. and Frans de Waal (2000). Any animal whatever: Darwinian building blocks of morality in monkeys and apes. *Journal of Consciousness Studies*, 7(1–2): 1–29.
- Fodor, Jerry A. (1975). The Language of Thought. Thomas Y. Crowell, New York.
- Fodor, Jerry A. (1987). Psychosemantics. The MIT Press, Cambridge, MA.
- Fodor, Jerry A. (1998). There are no recognitional concepts—not even red, part 2: the plot thickens. In In Critical Condition: Polemical Essays on Cognitive Science and the Philosophy of Mind, pages 49–62. The MIT Press, Cambridge, MA.
- Fodor, Jerry A. (2008). LOT 2: The Language of Thought Revisited. Oxford University Press, Oxford.
- Frege, Friedrich Ludwig Gottlob (1923). Logische Untersuchungen. Dritter Teil:
 Gedankengefüge. Beiträge zur Philosophie des deutschen Idealismus, III: 36– 51. Reprinted in Stoothoff, R. H. (Trans.) (1963). "Compound Thoughts" Mind

72(285): 1-17.

- Gabriel, Iason (2020). Artificial intelligence, values, and alignment. *Minds and Machines*, 30: 411–437.
- Gamble, Clive (2013). Settling the Earth: The Archaeology of Deep Human History. Cambridge University Press, Cambridge.
- Garde, Murray (2009). The language of fire: Seasonality, resources and landscape burning on the arnhem land plateau. In Russell-Smith, J., P. Whitehead, and P. Cooke, editors, *Culture, Ecology and Economy of Fire Management in North Australian Savannas: Rekindling the WURRK Tradition*, pages 85–164. CSIRO Publishing, Collingwood.
- Gardner, Howard (2011). Frames of Mind: The Theory of Multiple Intelligences. Basic Books, New York.
- Gibbard, Allan (1999). Morality as consistency in living: Korsgaard's kantian lectures. *Ethics*, 110(1): 140–164.
- Gilligan, Carol (1982). In a Different Voice. Harvard University Press, Cambridge, MA.
- Goebel, Randy (2008). Folk reducibility and AI-complete problems. In Dengel, A. R., K. Berns, T. M. Breuel, F. Bomarius, and T. R. Roth-Berghofer, editors, *KI 2008: Advances in Artificial Intelligence*, volume 5243 of *Lecture Notes in Computer Science*, pages 54–57. Springer, Berlin, Heidelberg. doi: 10.1007/978-3-540-85845-4_1.
- Goldberg, David E. (1987). Simple genetic algorithms and the minimal deceptive problem. In Davis, Lawrence D., editor, Genetic Algorithms and Simulated Annealing (Research Notes in Artificial Intelligence), pages 74–88. Morgan Kaufmann Publishers, Burlington, MA.
- Goldman, Alvin (1970). A Theory of Human Action. Prentice-Hall, Englewood Cliffs, NJ.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. The MIT Press, Cambridge, MA. http://www.deeplearningbook.org.
- Gowlett, John A. J. (2006). The Early Settlement of Northern Europe: Fire History in the Context of Climate Change and the Social Brain. *Comptes Rendus Palevol*, 5(1–2): 299–310.
- Gowlett, John A. J. (2016). The Discovery of Fire by Humans: A Long and Convoluted Process. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1696): 1–12.
- Gowlett, John A. J. and Richard W. Wrangham (2013). Earliest Fire in Africa: Towards the Convergence of Archaeological Evidence and the Cooking Hypothesis. *Azania: Archaeological Research in Africa*, 48(1): 5–30.
- Green, Steven and Peter Marler (1979). The analysis of animal communication. In Marler, P. and G. Vandenbergh, editors, *Handbook of Behavioral Neurobiology*, volume 3: Social Behavior and Communication, page 73–158. Plenum, New York.
- Greene, Joshua D. (2007). The Biology of Morality: Neuroscientists Respond to Killen and Smetana. *Human Development*, pages 1–2.
- Grossman, Sanford J. and Oliver D. Hart (1986). The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy*, 94: 691–719.
- Haddock, Steven H. D., Mark A. Moline, and James F. Case (2010). Bioluminescence in the Sea. Annual Review of Marine Science, 2: 443–493.

- Hadfield-Menell, Dylan and Gillian K. Hadfield (2019). Incomplete contracting and ai alignment. In Conitzer, Vincent, Gillian Hadfield, and Shannon Vallor, editors, AIES '19: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, pages 417–422. Association for Computing Machinery, New York.
- Hailman, J., M. Ficken, and R. Ficken (1985). The 'Chick-a-dee' calls of Parus atricapillus. Semiotica, 56: 191–224.
- Harris, Paul L. (2005). Conversation, pretense and theory of mind. In Astington,J. W. and J. A. Baird, editors, Why language matters for theory of mind, pages 70–83. Oxford University Press, Oxford.
- Hart, Oliver (1995). Firms, Contracts, and Financial Structure. Oxford University Press, Oxford.
- Haugeland, John (1985). Artificial Intelligence: The Very Idea. The MIT Press, Cambridge.
- Hauser, Marc D. (1996). The Evolution of Communication. MIT Press, Cambridge, MA.
- Hauser, Marc D., Noam Chomsky, and W. Tecumseh Fitch (2002). The Faculty of Language: What Is It, Who Has It, and How Did It Evolve? *Science*, 298: 1569– 1579.
- Hauser, Marc D. and W. Tecumseh Fitch (2003). What are the uniquely human components of the language faculty? In Christiansen, M. H. and S. Kirby, editors, *Language Evolution*, pages 158–181. Oxford University Press, Oxford.
- Hausman, Daniel M. (2012). Preference, Value, Choice, and Welfare. Cambridge University Press, Cambridge.
- Haynes, John (1895). Risk as an Economic Factor. *Quarterly Journal of Economics*, 9(4): 409–444.
- Henrich, Joseph (2016). The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species and Making Us Smarter. Princeton University Press, Princeton.
- Hou, Jianwei, Ann Kuzma, and John Kuzma (2009). Winner's Curse or Adverse Selection in Online Auctions: The Role of Quality Uncertainty and Information Disclosure. *Journal of Electronic Commerce Research*, 10(3): 144–154.
- Houck, Lynne D. (2009). Pheromone Communication in Amphibians and Reptiles. Annual Review of Physiology, 71: 161–176.
- Hubinger, Evan, Chris van Merwijk, Vladimir Mikulik, Joar Skalse, and Scott Garrabrant (2021). Risks from learned optimization in advanced machine learning systems. arXiv preprint arXiv:1906.01820, pages 1–39. https://arxiv.org/ abs/1906.01820.
- Hurford, James R. (2012). Language in the Light of Evolution II: The Origins of Grammar. Oxford University Press, Oxford.
- Hutto, Daniel D. (2012). Folk psychological narratives. The MIT Press, Cambridge, MA.
- Jackendoff, Ray (1997). The Architecture of the Language Faculty. The MIT Press, Cambridge, MA.
- Jackendoff, Ray (1999). Possible stages in the evolution of the language faculty. Trends in Cognitive Sciences, 3: 272–279.
- Jackendoff, Ray (2003). Foundations of Language: Brain, Meaning, Grammar, Evolution. Oxford University Press, Oxford.

- Jackendoff, Ray (2012). Your theory of language evolution depends on your theory of language. In Larson, Richard K., Viviane Déprez, and Hiroko Yamakido, editors, *The Evolution of Human Language: Biolinguistic Perspectives*, pages 63– 72. Cambridge University Press, Cambridge.
- Jackendoff, Ray and Steven Pinker (2005). The nature of the language faculty and its implications for evolution of language (reply to Fitch, Hauser, and Chomsky). *Cognition*, 97(2): 211–225.
- Jackson, Ryan Blake (2019). Toward morally sensitive robotic communication. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pages 715–717. IEEE. doi: 10.1109/HRI.2019.8673209.
- Jackson, Ryan Blake and Tom Williams (2019). Language-capable robots may inadvertently weaken human moral norms. In 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pages 401–410.
- Janssen, Theo M. V. (1997). Compositionality. In van Benthem, J. and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 417–473. Elsevier, Amsterdam.
- Janssen, Theo M. V. (2012). Compositionality: Its Historic Context. In Hinzen, Wolfram, Edouard Machery, and Markus Werning, editors, *The Oxford Handbook* of Compositionality, pages 19–46. Oxford University Press, Oxford.
- Jensen, Michael C. and William H. Meckling (1976). Theory of the firm: Managerial behaviour, agency costs and ownership structure. *Journal of Financial Economics*, 3(4): 305–360.
- Kamp, Hans and Barbara Partee (1995). Prototype Theory and Compositionality. Cognition, 57: 129–191.
- Kasirzadeh, Atoosa (2021). Reasons, values, stakeholders: A philosophical framework for explainable artificial intelligence. *arXiv preprint*, 2103.00752: 1–39. https://arxiv.org/abs/2103.00752.
- Kitcher, Philip (2006a). Between fragile altruism and morality: Evolution and the emergence of normative guidance. In Boniolo, Giovanni and Gabriele De Anna, editors, *Evolutionary Ethics and Contemporary Biology*, pages 159–77. Oxford University Press, Oxford.
- Kitcher, Philip (2006b). Biology and Ethics. In Copp. David, editor, The Oxford Handbook of Ethical Theory, pages 163–185. Oxford University Press, Oxford.
- Kitcher, Philip (2011). The Ethical Project. Harvard University Press, Cambridge, MA.
- Knight, Frank H. (1921). Risk, Uncertainty and Profit. University of Chicago Press, Chicago.
- Kohlberg, Lawrence (1958). The development of modes of moral thinking and choice in the years ten to sixteen. PhD thesis, University of Chicago, Irvine.
- Kohlberg, Lawrence (1963). The development of children's orientation toward a moral order: 1. sequence in the development of moral thought. Vita Humana, 6: 11–33.
- Kohlberg, Lawrence (1964). The development of moral character and ideology. In Hoffman, M. L., editor, *Review of child development research*, pages 383–431. Russell Sage Foundation, New York.
- Kohlberg, Lawrence (1969). Stage and sequence: The cognitive-developmental approach to socialization. In Goslin, D., editor, *Handbook of socialization theory and research*, pages 347–480. Rand McNally, Chicago.

- Kohlberg, Lawrence (1981). Essays on moral development, vol. I: The philosophy of moral development. Harper & Row, San Francisco.
- Kohlberg, Lawrence (1984). Essays on moral development, vol. II: The psychology of moral development. Harper & Row, San Francisco.
- Korsgaard, Christine M. (1996). *The Sources of Normativity*. Cambridge University Press, Cambridge.
- Korsgaard, Christine M. (2018). Fellow Creatures: Our Obligations to the Other Animals. Oxford University Press, Oxford.
- Koszegi, Botond and Matthew Rabin (2007). Mistakes in Choice-Based Welfare Analysis. American Economic Review, 97(2): 477–481.
- Krakovna, Victoria, Jonathan Uesato, Vladimir Mikulik, Matthew Rahtz, Tom Everitt, Ramana Kumar, Zac Kenton, Jan Leike, and Shane Legg (2021). Specification gaming: the flip side of ai ingenuity. https://deepmind.com/blog/ article/Specification-gaming-the-flip-side-of-AI-ingenuity.
- Kripke, Saul A. (1982). Wittgenstein on Rules and Private Language. Harvard University Press, Cambridge, MA.
- LaCroix, Travis (2020). Complex Signals: Reflexivity, Hierarchical Structure, and Modular Composition. PhD thesis, University of California, Irvine.
- LaCroix, Travis (2021). Reflexivity, Functional Reference, and Modularity: Alternative Targets for Language Origins. *Philosophy of Science*, 88(5): 1234–1245.
- LaCroix, Travis (2022a). Information and meaning in the evolution of compositional signals. *Philosophy of Science Association*, *Biennial Meeting 2022*. http:// philsci-archive.pitt.edu/20716/.
- LaCroix, Travis (2022b). Moral Dilemmas for Moral Machines. AI and Ethics. Forthcoming. https://doi.org/10.1007/s43681-022-00134-y.
- LaCroix, Travis (2022c). Using Logic to Evolve More Logic: Composing Logical Operators via Self-assembly. British Journal for the Philosophy of Science, 73(2): 407–437.
- LaCroix, Travis and Yoshua Bengio (2019). Learning from learning machines: Optimisation, rules, and social norms. arXiv preprint arXiv:2001.00006. https: //arxiv.org/abs/2001.00006.
- LaCroix, Travis and Alexandra Sasha Luccioni (2022). Metaethical Perspectives on 'Benchmarking' AI Ethics. arXiv preprint, 2204.05151: 1–39. https://arxiv. org/abs/2204.05151.
- LaCroix, Travis and Cailin O'Connor (2021). Power by Association. *Ergo*. Forthcoming.
- Laffont, Jean-Jacques (2000). *Incentives and Political Economy*. Oxford University Press, Oxford.
- Laffont, Jean-Jacques and David Martimort (2002). The Theory of Incentives: The Principal-Agent Model. Princeton University Press, Princeton.
- Legg, Shane and Marcus Hutter (2007). Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4): 391–444.
- Lehman, Joel and Kenneth O. Stanley (2008). Exploiting open-endedness to solve problems through the search for novelty. In *Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI)*, pages 329–336, Cambridge, MA. The MIT Press.
- Lishak, Robert S. (1984). Alarm Vocalizations of Adult Gray Squirrels. Journal of Mammalogy, 65: 681–684.

- Lugli, Marco, Hong Y. Yan, and Michael L. Fine (2003). Acoustic Communication in Two Freshwater Gobies: the Relationship Between Ambient Noise, Hearing Thresholds and Sound Spectrum. *Journal of Comparative Physiology A*, 189: 309– 320.
- Lust, Barbara (2006). *Child Language: Acquisition and Growth*. Cambridge University Press, Cambridge.
- Malle, Bertram F. (2015). Moral competence in robots? In Seibt, Johanna, Raul Hakli, and Marco Nørskov, editors, Sociable Robots and the Future of Social Relations, pages 189–198. IOS Press.
- Malle, Bertram F. (2016). Integrating robot ethics and machine morality: the study and design of moral competence in robots. *Ethics and Information Technology*, 18: 243–256.
- Malle, Bertram F. and Matthias Scheutz (2014). Moral competence in social robots. In 2014 IEEE International Symposium on Ethics in Science, Technology and Engineering, pages 1–6. doi: 10.1109/ETHICS.2014.6893446.
- Marco F. H. Schmidt and Hannes Rakoczy and Michael Tomasello (2012). Young children enforce social norms selectively depending on the violator's group affiliation. *Cognition*, 124(3): 325–333.
- Marcus, Gary and Ernest Davis (2020). Gpt-3, bloviator: Openai's language generator has no idea what it's talking about. MIT Technology Review. https://www.technologyreview.com/2020/08/22/1007539/ gpt3-openai-language-generator-artificial-intelligence-ai-opinion/.
- Markant, Douglas B. and Todd M. Gureckis (2014). Is it better to select or to receive? learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, 143(1): 94–122.
- Marler, Peter and Hans Slabbekoorn (2004). *Nature's Music*. Elsevier Academic Press, Amsterdam.
- Marschak, Jacob and Roy Radner (1972). *Economic Theory of Teams*. Yale University Press, New Haven.
- Mäthger, Lydia M., Eric J. Denton, N. Justin Marshall, and Roger T. Hanlon (2009). Mechanisms and Behavioural Functions of Structural Coloration in Cephalopods. *Journal of the Royal Society Interface*, 6: S149–S163.
- Maynard Smith, John and David Harper (2003). *Animal Signals*. Oxford University Press, Oxford.
- McCoy, Tom, Ellie Pavlick, and Tal Linzen (2019). Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In Korhonen, Anna and Lluís Màrquez David Traum, editors, *Proceedings of the 57th Annual Meeting* of the Association for Computational Linguistics, pages 3428–3448, Florence. Association for Computational Linguistics.
- McDermott, Drew (1987). A critique of pure reason. *Computational Intelligence*, 3: 151–160.
- Mehler, Jacques, Marina Nespor, Mohinish Shukla, and Marcela Peña (2006). Why is Language Unique to Humans? *Novartis Foundation Symposium*, 270: 251–280.
- Meir, Irit, Wendy Sandler, Carol Padden, and Mark Aronoff (2010). Emerging sign languages. In Marschark, M., editor, Oxford Handbook of Deaf Studies, Language, and Education, pages 267–280. Oxford University Press, Oxford.
- Mele, Alfred R. (1992). Springs of Action: Understanding Intentional Behavior. Oxford University Press, Oxford.

Mele, Alfred R. (2003). Motivation and Agency. Oxford University Press, Oxford.

- Miceli, Milagros, Julian Posada, and Tianling Yang (2022). Studying up machine learning data: Why talk about bias when we mean power? Proc. ACM Hum.-Comput. Interact., 6(GROUP): 1–14.
- Milligan, Karen, Janet Wilde Astington, and Lisa Ain Dack (2007). Language and theory of mind: meta-analysis of the relation between language ability and false-belief understanding. *Child Development*, 78(2): 622–646.
- Mitchell, Melanie, Stephanie Forrest, and John H. Holland (1992). The royal road for genetic algorithms: Fitness landscapes and GA performance. In Varela, F. J. and P. Bourgine, editors, *Proceedings of the First European Conference on Artificial Life*, pages 1–11. The MIT Press, Cambridge, MA.
- Nelson, Katherine (2005). Language pathways into the community of minds. In Astington, J. W. and J. A. Baird, editors, Why language matters for theory of mind, pages 26–49. Oxford University Press, Oxford.
- Ng, Andrew Y. and Stuart Russell (2000). Algorithms for inverse reinforcement learning. In Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000), pages 663–670. Association for Computing Machinery, New York.
- Normore, Calvin (1990). Ockham on mental language. In Smith, J-C., editor, *The Historical Foundations of Cognitive Science*, pages 53–70. Kluwer, Dordrecht.
- Ofek, Haim (2001). Second Nature: Economic Origins of Human Evolution. Cambridge University Press, Cambridge.
- Orseau, Laurent and Stuart Armstrong (2016). Safely interruptible agents. https://deepmind.com/research/publications/2019/ safely-interruptible-agents.
- Otte, Daniel (1974). Effects and Functions in the Evolution of Signaling Systems. Annual Reviews of Ecology and Systematics, 5: 385–417.
- Pagin, P. and D. Westerståhl (2010a). Compositionality I: Definitions and Variants. *Philosophy Compass*, 5(3): 250–264.
- Pagin, P. and D. Westerståhl (2010b). Compositionality II: Arguments and Problems. *Philosophy Compass*, 5(3): 265–282.
- Pandia, Lalchand, Yan Cong, and Allyson Ettinger (2021). Pragmatic competence of pre-trained language models through the lens of discourse connectives. arXiv preprint, 2109.12951: 1–13. https://arxiv.org/abs/2109.12951.
- Parker, Anna R. (2006). Evolving the narrow language faculty: Was recursion the pivotal step? In Cangelosi, Angelo, Andrew D. M. Smith, and Kenny Smith, editors, *The Evolution of Language: Proceedings of the 6th International Conference* (EVOLANG6), pages 239–246. World Scientific, Hackensack, NJ.
- Partee, Barbara Hall (1984). Compositionality. In Landman, F. and F. Veltman, editors, Varieties of Formal Semantics, page 281–311. Foris, Dordrecht.
- Pauer-Studer, Herlinde (2003). Christine m. korsgaard: Internalism and the sources of normativity. In Pauer-Studer, Herlinde, editor, Constructions of Practical Reason: Interviews on Moral and Political Philosophy, pages 50–69. Stanford University Press, Stanford.
- Peterson, Candida C. and Michael Siegal (1999). Representing inner worlds: Theory of mind in autistic, deaf and normal- hearing children. *Psychological Science*, 10(2): 126–129.

- Piaget, Jean (1959/1923). The language and thought of the child. Routledge & Kegan Paul, London.
- Piaget, Jean (1965/1932). The Moral Judgment of the Child. Kegan, Paul, Trench, Trubner & Co., London.
- Piaget, Jean (1969/1924). Judgment and reasoning in the child. Routledge & Kegan Paul, London.
- Pickering, Travis Rayne (2013). Rough and Tumble: Aggression, Hunting, and Human Evolution. University of California Press, Los Angles.
- Pinker, Steven and Ray Jackendoff (2005). The faculty of language: What's special about it? Cognition, 95(2): 201–236.
- Planer, Ronald J. and Kim Sterelny (2021). From Signal to Symbol: The Evolution of Language. The MIT Press, Cambridge, MA.
- Poulshock, Joseph W. (2006). Language and Morality: Evolution, Altruism, and Linguistic Moral Mechanisms. PhD thesis, University of Edinburgh.
- Reed, Russell D. and Robert J. Marks, II (1999). Neural Smithing: Supervised Learning in Feedforward Artificial Neural Networks. The MIT Press, Cambridge, MA.
- Rescorla, Michael (2019). The language of thought hypothesis. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2019 edition. https://plato.stanford.edu/ archives/sum2019/entries/language-thought/.
- Roth-Hanania, Ronit, Maayan Davidov, and Carolyn Zahn-Waxler (1976). Empathy development from 8 to 16 months: Early signs of concern for others. *Infant Behavior and Development*, 34(3): 447–458.
- Rothschild, Michael and Joseph Stiglitz (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics*, 93(4): 541–562.
- Rowlands, Mark (2012). Can Animals Be Moral? Oxford University Press, Oxford.
- Russell, Stuart (2019). Human Compatible: Artificial Intelligence and the Problem of Control. Viking, New York.
- Russell, Stuart and Peter Norvig (2021). Artificial Intelligence: A Modern Approach. Pearson, Hoboken, NJ, 4 edition.
- Sahlgren, Magnus and Fredrik Carlsson (2021). The singleton fallacy: Why current critiques of language models miss the point. Frontiers in Artificial Intelligence, 4: 682578.
- Samuelson, Paul A. (1938a). A note on the pure theory of consumer's behaviour. Economica, 5(17): 61–71.
- Samuelson, Paul A. (1938b). A note on the pure theory of consumer's behaviour: An addendum. *Economica*, 5(19): 353–354.
- Sappington, David E. M. (1991). Incentives in principal-agent relationships. Journal of Economic Perspectives, 5(2): 45–66.
- Savage, Neil (2019). How ai and neuroscience drive each other forwards. Nature, 571: S15–S17.
- Schlenker, Philippe, Emmanuel Chemla, Anne M. Schel, James Fuller, Jean-Pierre Gautier, Jeremy Kuhn, Dunja Veselinović, Kate Arnold, Cristiane Cäsar, Sumir Keenan, Alban Lemasson, Karim Ouattara, Robin Ryder, and Klaus Zuberbühler (2016). Formal Monkey Linguistics: The Debate. *Theoretical Linguistics*, 42(1– 2): 173–201.

- Schlosser, Markus (2019). Agency. In Zalta, Edward N., editor, The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, winter 2019 edition. https://plato.stanford.edu/archives/win2019/entries/agency/.
- Scott-Phillips, Thomas C. and Richard A. Blythe (2013). Why is Combinatorial Communication Rare in the Natural World, and Why is Language an Exception to this Trend? *Journal of the Royal Society Interface*, 10(88): 1–7.
- Sen, Amartya K. (1973). Behaviour and the concept of preference. *Economica*, 40(159): 241–259.
- Sen, Amartya K. (1977). Rational fools: A critique of the behavioral foundations of economic theory. *Philosophy & Public Affairs*, 6(4): 317–344.
- Sen, Amartya K. (1993). Internal consistency of choice. *Econometrica*, 61(3): 495– 521.
- Sen, Amartya K. (1997). Maximization and the act of choice. *Econometrica*, 65(4): 495–521.
- Sen, Amartya K. (2002). Rationality and Freedom. Harvard University Press, Cambridge, MA.
- Seyfarth, Robert M., Dorothy L. Cheney, and Peter Marler (1980a). Monkey Responses to Three Different Alarm Calls: Evidence of Predator Classification and Semantic Communication. *Science*, 210: 801–803.
- Seyfarth, Robert M., Dorothy L. Cheney, and Peter Marler (1980b). Vervet Monkey Alarm Calls: Semantic Communication in a Free-Ranging Primate. Animal Behaviour, 28(4): 1070–1094.
- Shah, Sunit N. (2014). Literature review: The principal agent problem in finance. The CFA Institute Research Foundation, pages 1–55.
- Shahaf, Dafna and Eyal Amir (2007). Towards a Theory of AI Completeness. In Amir, Eyal, Vladimir Lifschitz, and Rob Miller, editors, AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning, pages 150–155, New York. AAAI Press.
- Shapley, Lloyd S. (1953). Stochastic games. Proceedings of the National Academy of Sciences of the United States of America, 39(10): 1095–1100.
- Sinha, Koustuv, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela (2021). Masked language modeling and the distributional hypothesis: Order word matters pre-training for little. arXiv preprint, 2104.06644: 1–26. https://arxiv.org/abs/2104.06644.
- Sipper, Moshe, Ryan J. Urbanowicz, and Jason H. Moore (2018). To know the objective is not (necessarily) to know the objective function. *BioData Mining*, 11(21): 1–3.
- Smetana, Judith G. and Judith L. Braeges (1990). The Development of Toddlers' Moral and Conventional Judgments. *Merrill-Palmer Quarterly*, 36(3): 329–346.
- Smetana, Judith G., Wendy M. Rote, Marc Jambon, Marina Tasopoulos-Chan, Myriam Villalobos, and Jessamy Comer (2012). Developmental changes and individual differences in young children's moral judgments. *Child Development*, 83(2): 683–696.
- Smith, W. John (1997). The behavior of communicating, after twenty years. In Owings, D. H., M. D. Beecher, and N. S. Thompson, editors, *Perspectives in Ethology*, volume 12: Communication, pages 7–54. Plenum, New York.

- Soares, Nate, Benja Fallenstein, Eliezer Yudkowsky, and Stuart Armstrong (2015). Corrigibility. In AAAI Workshops: Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence. AAAI, New York, NY.
- Sober, Elliott and David Sloan Wilson (1998). Unto Others: The Evolution and Psychology of Unselfish Behavior. Harvard University Press, Cambridge, MA.
- Spence, Michael (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3): 355–374.
- Spence, Michael (1974). Market Signalling: Informational Transfer in Hiring and Related Processes. Harvard University Press, Cambridge, MA.
- Spitzner, M. J. E. (1788). Ausführliche beschreibung der korbbienenzucht im sächsischen churkreise, ihrer dauer und ihres nutzens, ohne künstliche vermehrung nach den gründen der naturgeschichte und nach eigener langer erfahrung. Leipzig.
- Steinert-Threlkeld, Shane (2016). Compositional Signaling in a Complex World. Journal of Logic, Language and Information, 25: 379–397.
- Steinert-Threlkeld, Shane (2020). Toward the Emergence of Nontrivial Compositionality. *Philosophy of Science*, 87(5): 897–909.
- Sterelny, Kim (2012). Language, gesture, skill: The co-evolutionary foundations of language. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367: 2141–2151.
- Sterelny, Kim (2014). A paleolithic reciprocation crisis: Symbols, signals, and norms. *Biological Theory*, 9(1): 65–77.
- Sterelny, Kim (2021). The Pleistocene Social Contract: Culture and Cooperation in the Hominin Lineage. Oxford University Press, Oxford.
- Sumers, Theodore R., Robert D. Hawkins, Mark K. Ho, Thomas L. Griffiths, and Dylan Hadfield-Menell (2022). Linguistic communication as (inverse) reward design. arXiv preprint, 2204.05091: 1–6. https://arxiv.org/abs/2204.05091.
- Sutton, Richard S. and Andrew G. Barto (2018). *Reinforcement Learning: An Introduction*. The MIT press, Cambridge, MA.
- Szabó, Zoltán Gendler (2012). The Case for Compositionality. In Hinzen, Wolfram, Edouard Machery, and Markus Werning, editors, *The Oxford Handbook of Compositionality*, pages 64–80. Oxford University Press, Oxford.
- Szabó, Zoltán Gendler (2020). Compositionality. In Zalta, Edward N., editor, *The Stanford Encyclopedia of Philosophy*. The Metaphysics Research Lab, Stanford University, Fall 2020 edition. https://plato.stanford.edu/archives/ fall2020/entries/compositionality/.
- Tallerman, Maggie (2007). Did our ancestors speak a holistic protolanguage? Lingua, 117: 579–604.
- Tallerman, Maggie (2012). Protolanguage. In Tallerman, Maggie and Kathleen R. Gibson, editors, *The Oxford Handbook of Language Evolution*, pages 479–491. Oxford University Press, Oxford.
- Thiel, Martin and Thomas Breithaupt (2011). Chemical Communication in Crustaceans: Research Challenges for the Twenty-First Century. In Breithaupt, T. and M. Thiel, editors, *Chemical Communication in Crustaceans*, pages 3–22. Springer, New York.
- Thoppilan, Romal, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo

Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le (2022). Lamda: Language models for dialog applications. *arXiv preprint*, 2201.08239: 1–47. https://arxiv.org/abs/2201. 08239.

- Tolmeijer, Suzanne, Markus Kneer, Cristina Sarasua, Markus Christen, and Abraham Bernstein (2020). Implementations in machine ethics: A survey. arXiv preprint, 2001.07573: 1–38. https://arxiv.org/abs/2001.07573.
- Tomasello, Michael (2008). Origins of Human Communication. The MIT Press, Cambridge, MA.
- Tomasello, Michael (2014). A Natural History of Human Thinking. Harvard University Press, Cambridge, MA.
- Tomasello, Michael (2018). The normative turn in early moral development. Human Development, 61(4–5): 248–263.
- Tomasev, Nenad, Kevin R. McKee, Jackie Kay, and Shakir Mohamed (2021). Fairness for unobserved characteristics: Insights from technological impacts on queer communities. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, page 254–265, New York, NY, USA. Association for Computing Machinery. https://doi.org/10.1145/3461702.3462540.
- Traxler, Matthew J., Megan Boudewyn, and Jessica Loudermilk (2012). What's special about human language? the contents of the 'narrow language faculty' revisited. *Language and Linguistics Compass*, 6(10): 611–621.
- Twomey, Terrence (2013). The Cognitive Implications of Controlled Fire Use by Early Humans. *Cambridge Archaeological Journal*, 23(1): 113–128.
- Twomey, Terrence (2014). How Domesticating Fire Facilitated the Evolution of Human Cooperation. *Biology and Philosophy*, 29(1): 89–99.
- Vaughan, Emmett J. (1997). Risk Management. Wiley, New York.
- Veres, Csaba (2022). Large language models are not models of natural language: they are corpus models. *arXiv preprint*, 2112.07055: 1–12. https://arxiv.org/abs/2112.07055.
- Vincent, Sarah, Rebecca Ring, and Kristin Andrews (2019). Normative practices of other animals. In Zimmerman, Aaron, Karen Jones, and Mark Timmons, editors, *Routledge Handbook of Moral Epistemology*, pages 57–83. Routledge, New York.
- von Frisch, Karl (1967). The Dance Language and Orientation of the Bees. Harvard University Press, Cambridge, MA.
- von Neumann, John and Oskar Morgenstern (1944). Theory of Games and Economic Behavior. Princeton University Press, Princeton.
- Wacewicz, Sławomir (2012). The narrow faculty of language: What is it, who has it, and how is it defined? *Theoria et Historia Scientiarum*, IX: 217–229.
- Warneken, Felix and Michael Tomasello (2006). Altruistic helping in human infants and young chimpanzees. Science, 31: 1301–1303.

- Warneken, Felix and Michael Tomasello (2007). Helping and cooperation at 14 months of age. *Infancy*, 11: 271–294.
- Wellman, Henry M. and Candida C. Peterson (2013). Theory of mind, development, and deafness. In Baron-Cohen, S., H. Tager-Flusberg, and M. V. Lombardo, editors, Understanding other minds: Perspectives from developmental social neuroscience, pages 51–71. Oxford University Press, Oxford.
- Wiener, Norbert (1960). Some moral and technical consequences of automation. Science, 131(3410): 1355–1358.
- Williams, Tom (2018). Toward ethical natural language generation for human-robot interaction. In Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, page 281–282. Association for Computing Machinery. doi 10.1145/3173386.3176975.
- Williamson, Oliver E. (1973). Markets and hierarchies: Some elementary considerations. The American Economic Review, 63(2): 316–325.
- Williamson, Oliver E. (1975). Markets and Hierarchies, Analysis and Antitrust Implications: A Study of the Economics of Internal Organization. The Free Press, New York.
- Wittgenstein, Ludwig (2009/1953). Philosophical Investigations. Wiley-Blackwell, Oxford, 4 edition.
- Wray, Alison (1998). Protolanguage as a Holistic System for Social Interaction. Language & Communication, 18: 47–67.
- Wray, Alison (2000). Holistic utterances in protolanguage: the link from primates to humans. In Knight, C., M. Studdert-Kennedy, and J. R. Hurford, editors, *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, pages 285–302. Cambridge University Press, Cambridge.
- Wray, Alison (2002). Dual processing in protolanguage: performance without competence. In Wray, A., editor, *The Transition to Language*, pages 113–137. Oxford University Press, Oxford.
- Wyatt, Tristram D. (2010). Pheromones and Signature Mixtures: Defining Species-Wide Signals and Variable Cues for Identity in Both Invertebrates and Vertebrates. Journal of Comparative Physiology A, 196: 685–700.
- Wynn, Karen and Paul Bloom (2014). The Moral Baby. In Killen, M. and J. G. Smetana, editors, *Handbook of Moral Development*, pages 435–453. Psychology Press, New York.
- Yudkowsky, Eliezer (2018). Ai alignment: Why it's hard, and where to start. https://intelligence.org/2016/12/28/ ai-alignment-why-its-hard-and-where-to-start/.
- Zahavi, Amotz (1987). The theory of signal selection and some of its implications. In Delfino, U. P., editor, *International Symposium on Biology and Evolution*, pages 294–327. Adriatica Editrice, Bari.
- Zahn-Waxler, C., M. Radke-Yarrow, E. Wagner, and M. Chapman (1992). Development of concern for others. *Developmental Psychology*, 28(1): 126–136.
- Zuberbühler, Klaus (2018). Combinatorial capacities in primates. Current Opinion in Behavioral Sciences, 21: 161–169.
- Zuberbühler, Klaus, Dorothy L. Cheney, and Robert M. Seyfarth (1999). Conceptual semantics in a non-human primate. *Journal of Comparative Psychology*, 113(1): 33–42.