Content Indeterminacy in Egan's Deflationary Account of Mental Representation

Steven Nathan McGannon
Master's thesis
Linguistic Diversity and Digital Humanities
Cognitive Science
Faculty of Arts
University of Helsinki
May 2022

HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

Tiedekunta – Fakultet – Faculty Faculty of Arts	Koulutusohjelma – Utbildningsprogram – Degree Programme Linguistic Diversity and Digital Humanities
Opintosuunta – Studieinriktning – Study Track Cognitive Science	
Tekijä – Författare – Author Steven Nathan McGannon	
Työn nimi – Arbetets titel – Title Content Indeterminacy in Egan's Deflationary Account of Mental Representation	

Työn laji – Arbetets art – Level MA	Aika – Datum – Month and year May 2022	Sivumäärä– Sidoantal – Number of pages 41

Tiivistelmä – Referat – Abstract

This thesis examines the notion of representational content in philosopher Frances Egan's deflationary account of mental representation (DAMR) for cognitive neuroscience. In particular, this thesis explores to what extent DAMR's pragmatic account of representational content answers to the content determinacy constraint. The content determinacy constraint is a constraint on an account of mental representation which holds that an account must provide the basis for attributing determinate content to the posited computational states. For example, the constraint requires that an account must specify the conditions under which a given representational content — such a *fly* — is attributed to a given representational vehicle — such as some particular structures in a frog's brain. A variety of naturalistic "tracking theories" — broadly, theories which attempt to answer to such a constraint in terms of a privileged naturalistic relation holding between some object and the inner mental state in question — are often held to have thus far failed to satisfactorily meet this constraint.

Egan's deflationary account promises to address this constraint in a way which departs from naturalistic tracking theories (and from several other theories as well). It promises to do this by proffering a notion of representational content which says that such content is fundamentally *pragmatic*, and that this ultimately allows the account to avoid the problems typically associated with tracking theories. I will examine how DAMR purports to meet the content determinacy constraint via its pragmatic account of content. To this end, I will raise a content indeterminacy concern for one of the pragmatic functions of content DAMR maintains. I will examine several potential lines of response to this indeterminacy concern, some of which I will evaluate as more or less successful.

The first chapter introduces the central research aims, methods, and an ethical research statement. The second chapter consists in the theoretical backdrop against which DAMR is situated, in particular, with regard to how the notions of "representation" and "computation" have been employed in cognitive science historically. The third chapter consists in a detailed explication of the components of DAMR, including a careful differentiation between its realist and computational elements and its pragmatic elements. This is followed by an explication of how DAMR purports to meet the content determinacy constraint. The fourth chapter consists in a theoretical concern which suggests some pragmatic elements of DAMR might be vulnerable to a content indeterminacy concern; and, a variety of potential responses to the indeterminacy concern raised. The fifth chapter consists in setting out the central conclusion of the thesis and suggesting four potential areas for future investigation.

The central conclusion of this thesis is that DAMR faces a "limited" content indeterminacy concern for at least one of its pragmatic elements. However, I highlight that there are several ambiguities which require resolution prior to consideration if the concern raised extends to other elements of the account. I propose that the results of this thesis demonstrate the need for further clarification with regard to the way in which DAMR is committed to the content determinacy constraint. Ultimately, the concerns raised along with the ambiguities noted should serve to facilitate further philosophical evaluation of DAMR.

Avainsanat – Nyckelord – Keywords
deflationary account
mental representation
content determinacy
Frances Egan

Säilytyspaikka – Förvaringställe – Where deposited

Muita tietoja – Övriga uppgifter – Additional information

Acknowledgements

I want to thank my thesis advisor Anna-Mari Rusanen (University of Helsinki) and Jesse Kuokkanen (University of Helsinki) for their guidance and feedback throughout the thesis process. Additionally, I want to thank Riikka Möttönen (University of Helsinki) for her helpful comments, and Otto Lappi (University of Helsinki) for his helpful correspondence. I want to further thank Gabriel Sandu (University of Helsinki) and Jonathan Ellis (University of California, Santa Cruz) for discussions early in the thesis process which contributed to the ultimate direction of this thesis. I also want to thank Eeva Karjalainen: without you, writing this thesis would not have been possible. Additionally, I want to thank my parents for all their support. And finally, I want to thank all my friends and colleagues, the Linguistic Diversity and Digital Humanities program at University of Helsinki, and University of Helsinki.

Table of Contents

1. Introduction

1.1. Research Aims and Methodology

In this thesis, I will examine the deflationary account of mental representation (DAMR) proffered by philosopher Frances Egan. DAMR aims to provide an account of the nature and function of representation in cognitive neuroscience. To this end, DAMR maintains a realist construal of mental representational vehicles and a pragmatic account of mental representational content (Egan, 2010, 2014, 2017, 2018, 2020). A realist construal of representational vehicles takes such vehicles to be actually-existing states or structures in a cognitive system. Meanwhile, a pragmatic account of representational content takes such content to be ascribed by theorists for broadly *pragmatic* (or heuristic) reasons (Egan, 2020). In particular, the account is motivated by the search to offer an explanation of representational content. Evaluating DAMR's explanation thereof will form the central investigation of my thesis. To this end, the primary research problem that I address in my thesis is: does DAMR meet the *content determinacy constraint* on an account of representational content? In brief, the content determinacy constraint says that an account of representational content must provide the basis for the attribution of *determinate content* to a posited state or structure (Egan, 2020). I will narrow my focus to one of the pragmatic functions which DAMR maintains provides part of such a basis, namely, the "placeholder" pragmatic function of mental representational content (PF). I will raise a theoretical concern which states that PF might be vulnerable to content *indeterminacy*, and that this threatens to extend to other pragmatic aspects in DAMR. This will serve to address the primary research problem: does DAMR meet the *content determinacy constraint*; and, as I will show, will also open up avenues for further inquiry.

Broadly put, this thesis is situated in the context of cognitive science. Specifically, it is concerned primarily with theoretical topics in cognitive neuroscience. Thus, this thesis aims to address the relevant outstanding questions in cognitive neuroscience (the nature and function of mental representation, i.e.) in a philosophically-oriented way. The methodological approach is, thus, firmly rooted in the philosophical tradition: I engage with the philosophical and cognitive scientific theoretical literature and employ conceptual philosophical exegesis,

conceptual analysis, and argumentation. The aim of these broadly philosophical methods employed is the specification of an account which is concerned with cognitive neuroscience (in particular, *computational* cognitive neuroscience).

The layout of this thesis will be as follows. This thesis is divided into six chapters. Each chapter is further divided into subsections. This chapter (chapter 1) will consist in an introduction to the research aims of the thesis, its methods, and the central research problem addressed (section 1.1), followed by an explanation of the theoretical significance of the topic and the justification for undertaking this project (1.2), and finally, a statement on the ethics involved (1.3). Chapter 2 will consist in explicating the broader theoretical backdrop of the notion of representation and its place and use within cognitive science (2.1. - 2.2.), and an explication of the *content determinacy constraint* for theories of representational content. Chapter 3 will consist in explicating DAMR in full. 3.1 will consist in an explication of the "computational theory proper" in articulation with DAMR 3.2. will consist in an explication of the two kinds of content — mathematical and cognitive — which DAMR maintains. 3.3. will consist in an explication of DAMR's realist construal of representational vehicles. Finally, 3.4. will consist in an explication of how DAMR purports to answer to the content determinant constraint outlined in chapter 2, section 2 (3.4.). Chapter 4 will consist in evaluating whether or not DAMR provides the basis to answer the content determinacy constraint. 4.1. will consist in explicating previous criticisms to DAMR. In section 4.2., I will raise a theoretical concern which suggests that at least one of the pragmatic functions DAMR maintains might *not* provide the basis for the attribution of determinate content in some cases. 4.3. consists in raising five potential responses to the concern raised in 4.2. and addressing those responses in turn. Chapter 5 consists in presenting my central conclusion — which includes answering the central research problem raised in this thesis and, additionally, suggesting four considerations for potential future research. Chapter 6 comprises the references cited throughout the thesis.

1.2. Theoretical Significance and Justification

The theoretical significance of examining DAMR's commitment to content determinacy constraint — and to what extent it ultimately meets this constraint is as follows. First, the

notion of *representation* is a ubiquitous, central theoretical posit invoked broadly throughout cognitive science. It has often been referred to as *the* central explanatory posit of cognitive and a background theoretical assumption underlying much of cognitive scientific theorizing (Fodor, 1981; Ramsey, 2007; Egan, 2020; Rescorla, 2020; Smortchkova et al., 2020). Meanwhile, the *content determinacy constraint* (or various versions thereof) have proven to be a particularly elusive constraint to meet. To this end, a broad cross-section of theoretical-philosophical approaches have been advanced in attempt to address this issue. Thus far, no such approach has succeeded in providing a widely-accepted, satisfactory account. In chapter 2 (2.2.), for example, I will explicate a variety of naturalistic "tracking theories" which have failed to satisfactorily answer to this constraint. This backdrop against which DAMR is developed — a failure of a broad cross-section of theories to answer to the content determinacy constraint — points to the theoretical importance of developing an account which *does* so answer to this constraint. This forms part of the justification for this thesis.

Second, DAMR promises to provide a deflationary approach to answer the challenge just described, and indeed, to the challenge of characterizing the nature and function of representation in computational cognitive neuroscience. DAMR purports to answer the content determinacy constraint by demonstrating that a combination of a theorist's explanatory context along with additional pragmatic considerations ultimately serve to provide the basis for the attribution of determinate content to the states or structures posited in the computational theory proper (Egan, 2020). Clarification of how the account purports to do — and if it succeeds in doing so — is, thus, a philosophically and theoretically signifiant task. It is important because, *if* DAMR turns out to successfully meet the constraint that a broad variety of previous accounts of representation have failed to meet, *then* a central theoretical posit of cognitive neuroscience will have a robust, more justified theoretical grounding in DAMR. For example, DAMR could, if successful, suffice to answer a number of various outstanding theoretical questions surrounding the nature and function of mental representation: the *nature* of representational content (if content is, as DAMR suggests, fundamentally *pragmatic*), the status of computational theorizing in cognitive neuroscience more broadly, and to what extent competitor theories — such as naturalistic tracking theories — ought to continue to be theoretically pursued. The theoretical significance attached to what DAMR promises serves to provide a further justification for this thesis.

Finally, as a more minor point, previous critical evaluations of DAMR have tended to focus on the "essentiality" of the kinds of representational content DAMR maintains (see chapter 4.1. for discussion) (Ramsey, 2020; and see Coelho Mollo's (2017) discussion of various criticisms DAMR has faced). Such responses to DAMR have, thus, often not focused on how DAMR purports to meet the content determinacy constraint. Thus, as another point of significance, this thesis aims, in part, to contribute to evaluating this aspect of DAMR.

1.3. Statement on Research Ethics

This thesis is theoretical in nature, and I do not collect any personal data having to do with individuals for analysis. Thus, I claim no ethical concerns with regard to individual data collection or data preservation to claim. More broadly, I do not collect data or information about any individuals, nor do I discuss or mention any individuals (beyond any authors cited). Thus, I claim no ethical concerns with regard to people or animals. Additionally, my thesis also does not involve any discussion which would concern broader social notions — such as politics, nationality, immigration, gender, race, etc. — nor does it have any direct or obvious export to such discussions. Finally, I do not claim any conflicts of interest.

2. The Theoretical Backdrop and the Content Determinacy Constraint

In this chapter, I will outline the theoretical backdrop against which DAMR has been developed. In 2.1., I will first introduce the notion of representation within cognitive science. Second, I will highlight the role of representation in the computational-representational theory of mind (CRTM). Finally, I will explicate an important distinction between representational vehicles and representational content, which will serve as the subject matter for which DAMR aims to account. In 2.2., I will first introduce the content determinacy constraint on an account of representational content. Second, I will outline how a variety of naturalistic theories have aimed to account for this constraint, and I will highlight the problems they have faced in doing so. Finally, I will explicate and conclude with DAMR's theoretical motivations: the mismatch between the failure of naturalistic tracking theories to provide the basis for the

attribution of determinate content along with such contents continued attribution in cognitive neuroscientific practice, and Noam Chomsky's (1995) eliminativist account of representational content.

2.1. Representation in Cognitive Science

In this section, I will outline the theoretical backdrop against which DAMR has been developed. I will first outline the central role representation has been taken to play in cognitive science. Second, I will outline the *computational-representational theory of mind* (CRTM) and note where DAMR fits into this theoretical framework. Finally, I will explicate the *vehicle-content distinction*, which will ultimately serve to help make clear the explication of DAMR to follow in chapter 3.

The nature and function of representation is a significant, outstanding question in cognitive science. The notion of representation has played a central role throughout the history of cognitive science (Smortchkova et al., 2020). It is often taken to be an assumption of cognitive scientific practice and has been called the single most important explanatory posit of cognitive science generally (Fodor 1981; Shea, 2018; Rescorla, 2020). The cognitive revolution in the 1950s saw the notion of inner mental representations playing a central role in the then-burgeoning field of cognitive science. Stemming from its direct reaction to behaviorism, cognitive science has generally been described as assuming higher cognitive capacities involve inner mental representations (Miller, 2003; Ramsey, 2007).[1] Mental representation has been invoked in explanations ranging across psychological abilities, including perception, memory, reasoning, action (Shea, 2018).

Many projects have been carried out elucidating representation. Despite differences between various accounts, the notion of mental representation within cognitive science has traditionally been construed as involving a physical object that is instantiated in the brain and

[1] The cognitive revolution in the 1950s saw the development of mental representation playing this crucial role (Smortchkova et al., 2020). In large part, this grew out of a reaction to the predominant psychological paradigm behaviorism. In broad terms, behaviorism was a theoretical movement within psychology (and in particular, within psychology in the United States) which proposed to redefine psychology as the "science of behavior." Behaviorists argued that, since mental events are not publicly observable, objective psychological evidence must be behavioral (Miller, 2003). Though work by European psychologists and cognitivists like Miller and Broadbent in the 1950s started to move away from this picture, a seminal conference on artificial intelligence in 1956 is often taken as the point from which cognitive science as a discipline burgeoned (Bechtel and Herschbach, 2010). These developments are termed the "cognitive revolution", which thus initially grew largely as a reaction to behaviorism.

possesses semantic properties (Pitt, 2018; Smortchkova et al., 2020). Representations have also sometimes been employed in other ways. For example, connectionist models posit representations at the level of network node activation patterns (Smolensky, 1990).[2] It is thus often held that an aim (or even *the* aim) of cognitive science is to identify representations and determine how they are constructed and manipulated in the cognitive mechanisms that control behavior (Bechtel, 2016).

The Computational-Representational Theory of Mind

The central role representation has played as theoretical posit in cognitive science is perhaps best exemplified by what Smortchkova et al. (2020) call the *computational-representational theory of mind* (CRTM). Generally characterized, CRTM says: the mind, or mental processes, involve computations performed over representations. CRTM is a combination of two theses about the mind: the *representational theory of mind* (RTM) and the *computational theory of mind* (CTM). Generally characterized, CTM holds that the mind is *computational*; and, RTM holds that the mind is *representational*.[3] To better understand CRTM, I will explicate each constituent notion in turn, and then, together under CRTM.

RTM (sometimes, 'representationalism') is a framework about the mind. It has been characterized as the working assumption of cognitive science (Von Eckardt, 2012). Generally put, RTM holds that the mind is *essentially* representational (Smortchkova et al., 2020). In other words, the mind involves some kind of representational capacities. It is important to highlight RTM is not a *single, unified theory*. Rather, it is a general framework about how the mind works. That said, mental representations on RTM are often construed as physical things — or objects — with meaning or semantic properties (Shea, 2018).

Meanwhile, CTM (sometimes, 'computationalism') holds that vehicles which are manipulated in cognitive operations do not have semantic properties, only syntactic ones

[2] Additionally, other frameworks have questioned the need to invoke inner mental representations in cognitive science *at all* (Rupert, 2009). Such approaches include viewing the mind as an emergent activity of a larger environmental system, stressing the dynamic and coupled non-representational interaction of the organism and the environment. See, for example, dynamical systems approaches (Port and van Gelder 1995) and some of the so-called "4E" family of approaches: embodied/situated, embedded, enactive and extended cognition (for an early example Varela et al. (1991)).

[3] Importantly, RTM and CTM are sometimes used interchangeably or taken to be constitutively connected (Smortchkova et al., 2020). As Fodor (1981) famously phrased it, there is "no computation without representation".

(Smortchkova et al., 2020).[4] Importantly, "vehicles" here are *not representational vehicles*, since CTM need not employ representations. Rather, the vehicles in CTM are purely syntactic entities. Alan Turing's (1936) abstract mathematical model idealized computing device provided the backdrop for the development of the core of classical CTM (Rescorla, 2016).[5] On CTM, the mind is a similar, computational system resembling the central notions involved in Turing's idealized model of a computing device. Centrally, that the mind computes mental symbols. The nature of these symbols is variously conceived depending on the particular theory in question. Various versions of CTM have been held by many prominent philosophers and theorists (see for early examples: Putnam (1967), Fodor (1975), Newell and Simon (1976)).

Thus, the classical construal of CRTM holds to a conjunction of CTM and RTM . This conjunction holds, generally put: the mind or cognition is a computational process (*a la* CTM) which is defined over linguistically structured representations (*a la* RTM) (Fodor, 1975; Sprevak and Colombo, 2018; Smortchkova et al., 2020). One of the most classic examples of a CRTM is Jerry Fodor's (1975) "language of thought" (LOT) hypothesis. LOT claims that mental representation has a linguistic structure and that thought takes place in a mental language (which Fodor calls "*mentalese*"). According to LOT, CRTM is also committed to mental representations encoding propositional content in virtue of its syntax. Thus, there is a generally *computational* process over a variety of *representations* (here, a specifically *mental* "language").

The general theoretical appeal of the CRTM framework (and hence, theories like LOT), is that such a framework has offered promising solutions to various outstanding philosophical problems about the mind. For example, Smortchkova et al. (2018:4-5) identify CRTM as helpfully fitting mental states into a physicalist (or naturalistic) framework (since

---

[4] Sometimes, CTM is also called the computational theory of *cognition* (CTC) (Cummins, 1989; Piccinini, 2020). On this distinction, CTM is construed as a "strong" thesis about the mind which holds that both cognition *and consciousness* are computational, whereas CTC is a "weak" thesis about the mind which holds that at least *cognition*, but *not* necessarily *consciousness*, is computational. Careful attention to this distinction is not necessary for this thesis. This is because both CTM and CTC hold that *cognition* is computational, and DAMR — as I will explicate in chapter 3 — is centrally concerned with *cognitive* capacities (in computational cognitive neuroscience, i.e.). Closer attention to how DAMR treats *consciousness* is, thus, best left to future projects.

[5] Prior to Turing, many philosophers had conceived of the mind as involving *something like* or *related to* what would later be formulated as computation. For a helpful historical discussion see Isaac, 2018. However, both the notion of computation and of computers is usually traced back to Alan Turing's work on the foundations of mathematical computation and artificial intelligence. In this regard, Turing is generally credited with the first comprehensive abstract mathematical characterization of computation and of computers (Sprevak and Colombo, 2018; and see Turing, 1936, 1950).

computation can be *physically* implemented), as helping to explain how mental events could be *causal* events (since computational processes realized in neuronal structures as vehicles can allow for causal efficacy), and how to account for *reasoning* (via a computationally proceeding process, i.e.) — among others.

As I will explicate in the chapters to follow, DAMR will offer a novel take on CRTM. It will sit firmly against the backdrop of CTM — maintaining that, indeed, a computational "theory proper" is entirely sufficient to explain the success of a cognitive system at a given cognitive capacity. It will also take seriously the idea that representation plays an essential role in the explanatory practice of cognitive science (as in RTM), but, as I will explicate in chapter 3, it will construe representations as fundamentally *pragmatic* and, hence, *inessential* to the computational theory proper.

The Vehicle-Content Distinction

One of the central distinctions in discussions surrounding the nature and function of representations in cognitive science is the distinction between representational vehicles and representational content. I will now discuss theoretical distinction which will be crucial to understanding the account to be evaluated.

The vehicle-content distinction is a distinction about the elements of representation. It has been widely argued that the notion of representation *presupposes* a distinction between representational *vehicles* and representational *content* (Dennett 1991, Millikan 1991, 1993, Egan, 2014, 2018, 2020).[6] A representational *vehicle* is typically characterized as a physically realized state or structure that carries or bears content. Furthermore, representational vehicles are typically characterized as *causally explanatory* in a given cognitive process (Hurley, 1998; Hutto and Myin, 2020; Egan, 2020). Representational content is, then, *carried* by a vehicle just in case it represents things to be a certain way.

What physically realized states or structures can count as or serve as a representational vehicle might differ depending on the framework involved. For example, if the mind is assumed to be something like a classical von Neumann architecture, then the representational vehicles might be the data structures posited. Meanwhile, if the mind is assumed to be

[6] See also the earlier Dennett (1978:214) for an early, implicit example of the distinction.

something like a connectionist network, then the representational vehicles might be the network node activation patterns posited (Von Eckardt, 2012; Shea, 2007). Thus, there is no consensus across theoretical frameworks about exactly what kind of structures or states can or should count as representational vehicles.[7]

Meanwhile, a representation is also often said to have representational *content* (sometimes referred to as *intension*). The content (or intension) of a representation is typically considered to have to do with whatever things — objects, properties, qualities, etc. (sometimes referred to as *extensions*), variously construed — that a representation is said to *refer to* or be *about* (Ryder, 2004; Shea, 2018).[8] For example, the representational content of someone's memory of Helsinki Cathedral might be: *that is a very beautiful building*. Likewise, a map of Berkeley, California (a kind of representation, i.e.) is *about* Berkeley, California (the map's representational *content*, i.e.). Thus, content is often described or identified as a *semantic* property of representation (Cummins, 1989). In relation to representational vehicles, representational content is what a vehicle *carries* or *bears*. And as previously noted, it is generally the *vehicle* — not representational content itself — which is construed as causally efficacious. Furthermore, representations are also often said to have "satisfaction conditions."[9] The satisfaction condition of a given representation is the condition under which it represents *accurately* (Egan, 2020). For example, the satisfaction condition on an inner mental representation of a *cat*, say, might be construed as the condition of their being, in fact, a cat.[10]

---

[7] Ramsey (2007), for example, argues that connectionist network node pattern activation is not really representational *at all*, but rather, functions as "mere causal relays".

[8] Representational content will often be said to have *distal* content (typically, an object or property in the environment, e.g.) or *proximal* content (typically, something *not* in a cognitive subject's environment — such as a retinal image, e.g.). I will not distinguish closely between these kinds of content for the reason that, on DAMR, *either* kind can be an acceptable kind of content ascription (it will depend on the specific explanatory context and additional pragmatic considerations which, on DAMR, provide the basis for any content attribution (see chapter 3.4.)).

[9] Smortchkova et al. (2020:2) identify *satisfaction conditions* and *content* as two distinct semantic properties of representation.

[10] Sometimes, representations are said to act as a "surrogate" or a "stand in" for what they represent (Cummins, 1989; Grush, 1997; Piccinini, 2020; Egan, 2020). For example, a photograph of ice cream on an ice cream advertisement might be said to "stand in" for the ice cream so advertised. Similarly, a working thermostat is a representation of the current temperature because it stands in for that temperature, etc. Exactly how a representation might act in such a way is beyond the scope of this thesis. However, Egan ultimately does think DAMR can fulfill this sort of common construal of representational function, too — see Egan (2020) for her explication.

One of the central outstanding questions in theoretical discussions around cognitive science is how to specify the content of representations.[11] This is the basis of the *content determinacy constraint* on an account of representational content. I will dedicate the next subsection (2.2.) to a detailed statement of this constraint, as the central problem of this thesis will involve raising a theoretical concern involving how one of the pragmatic functions of content DAMR maintains meets this constraint (this will form the basis of chapter 4).

2.2. The Content Determinacy Constraint

In this section I will explicate the content determinacy constraint on an account of mental representational content. First, I will explicate what the constraint is supposed to be, paying attention to two different ways the constraint can be understood and clarifying which way I will understand it for this thesis. Second, I will look at several notable "tracking theories" and how they attempt to answer to the content determinacy constraint so understood. Finally, I will show how the various tracking theories' failures to provide the basis for content determinacy, along with Chomsky's (1995) eliminativist theory of representational content together from the two central theoretical motivations for DAMR.

Egan (2020) identifies four — as she characterizes them — "widely accepted" constraints on a theory of mental content. For the purposes of this thesis, I will limit my discussion to just one of these, namely, the content determinacy constraint.[12] Before I do so, it will be important to clarify how DAMR is committed to this constraint. It is central to understanding how DAMR purports to meet the constraint to follow that this constraint is generally applied to accounts of *mental* content. This is important because DAMR will make a distinction between two kinds of representational content: mathematical content and cognitive content. In DAMR, *cognitive* content is supposed to fill the role of what is usually considered to be *mental* content. Since these constraints are generally applied to theories of mental content, and cognitive content is supposed to fill that role, DAMR thus answers to the

[11] This includes, for example, extensive debates about the nature of representations and representational content — a variety of which I will explicate in the sections to follow.

[12] Other constraints identified by Egan include a commitment to *naturalism* (as Egan (2020) construes it: a non-intentional specification of at minimum a sufficient condition for a given content), *conformation to neuroscientific practice* (the account should make sense of how representations have been posited *in practice*, i.e.), and the possibility for *mis*representation (see the discussion to follow, and footnote 20).

constraint outlined below with reference to *cognitive* content and *not* mathematical content.[13] Ultimately, Egan will argue that DAMR does indeed meet this constraint.

Egan (2018, 2020) identifies the following (what she refers to as) widely-accepted constraint on an account of mental representational content:

> "The account should provide the basis for attributing *determinate contents* to computational states or structures," (Egan, 2020)

This constraint says that an account of representational content should answer to content determinacy (or sometimes content *determination*): it should be able to specify some set of conditions on which computational cognitive theorists attribute determinate content to the posited representational states or structures in their theories.

It will be important to disambiguate two ways of understanding this constraint as-is (as Egan formulates it above, i.e.). The first way of understanding this constraint is to understand it as saying that an account of representational content should provide the basis for attributing *a certain kind of representational content*, namely, *determinate* contents. The second way of understanding this constraint is to understand it as saying that such an account should provide the basis for "mere" *attribution* (of something *considered to be* "determinate content" among theorists, i.e.). This is an important distinction because the first understanding construes the constraint as being about providing the basis for attributing a certain kind of content (namely, *determinate* content), while the second understanding construes the constraint as being about providing the basis for a theorist's *attribution* (regardless of the nature of what they attribute).[14]

---

[13] Although Egan clearly argues for how she sees DAMR's construal of content as answering to the four constraints on mental content, nonetheless, there is room to ask if it is really correct to *exempt* the mathematical content of the computational theory proper from meeting these same constraints. It is outside the scope of this thesis to address this question.

[14] For textual support, note that Egan states, "…the account should provide the basis for *determinate* contents," (Egan, 2020:14). This seems to indicate that the account should provide such a basis for a certain kind of content, namely, *determinate* content (as opposed to *indeterminate* content; this is the first understanding outlined). However, Egan also states that DAMR "…does this by explicitly recognizing the role of explanatory interests and other pragmatic considerations in determining content ascription," (Egan, 2020:14). This latter statement indicates a determination of content *ascription or attribution* (the second understanding outlined). Taken together these statements (along with similar others — see for example Egan, 2020:5-6, e.g.) represent a seemingly ambiguous presentation of the content determinacy constraint in question. Resolving this ambiguity would have consequences for the worry presented in chapter 4 (4.2.) and for understanding DAMR more generally. It is outside the scope of this thesis to provide such a resolution, but I will suggest that further investigation is crucial in this regard.

Taking note of the preceding distinction is important for the following reason: in chapter 4 of this thesis, I will raise a theoretical concern for DAMR directly related to how it meets the content determinacy constraint. Importantly, I will raise this concern with regard to the *first* understanding of the content determinacy constraint outlined above (that it is foremost about the attribution of *determinate* content, e.g.). Thus, I will focus on this understanding of the content determinant constraint. Consequently, a central limitation of the results of my investigation is that they do not necessarily apply to the *second* understanding of the constraint outlined above (which is foremost about *attribution*, e.g.). However, I will return to the second understanding of the constraint again in chapter 4 (4.3.) when I evaluate the concern I raise in 4.2.; and, in chapter 5, where I will point to future research directions.

Given the above qualification I will now explicate what *determinate content* is supposed to be. Doing so will serve to clarify what kind of content is supposed to be attributed in a computational cognitive theory to meet the content determinacy constraint. Some theorists have characterized the content determinacy constraint as a central outstanding theoretical problem with regard to representational content involves: how exactly to specify what content a given representation has. On this understanding, content determinacy is about how to specify the content of a given representation. Generally put, *determinate* content is the specific, unique content of a given representation. A general way of asking what *determinate* content a given representation has is the following:

> In virtue of what does the content of a given mental representation have some content A and *not* some other content B (where A ≠ B)?

There are many different variations on the content determinacy constraint as a *problem* about how to specify what content a given mental representation has *rather than* some other content (how to specify that a representation's content is *A rather than B*, i.e.) (Fodor, 1990; Von Eckardt, 2012; Neander, 2017).[15][16] Versions of this problem have been taken to be a

[15] This problem is related to — or sometimes used interchangeably with — the so-called "disjunction problem" (Fodor, 1990). The disjunction problem is the problem which asks how a theory assigns unique mental content to mental states *rather than* some other, equally viable coextensive (*disjunctive*) content.

[16] Egan also addresses another variant of content determinacy concerns: *Quinean* indeterminacy (see Quine, 1960; and for Egan's discussion see Egan (2020:11)). It is beyond the scope of this thesis to address additional determinacy concerns, and I will thus limit my discussion of content determinacy to the constraint as outlined in this section (2.2).

foundational problem for cognitive science more generally (Von Eckhardt, 2012). It will suffice for the purposes of this thesis to take into account the general notion outlined above. This is because DAMR commits itself to a general, widely-accepted notion of the content determinacy constraint (Egan, 2020).[17]

To better understand the problem posed in this constraint, consider the following widely-employed example.[18] Frogs exhibit fly-catching behavior which is usually attributed to some inner state of the frog. A frog has some inner states or structures which serve as vehicles for some representational content when they exhibit tongue-lashing behavior. But what exactly does the frog's inner state *represent*? In other words, what exactly is the representational *content* of the frog's inner state? Numerous candidates have been proposed — and defended for various reasons — at one time or another by a broad cross-section of philosophers. For example, one might think the representational content is *fly* because it is a fly which triggers the frog's tongue-lashing. However, what the frog is seeking is not *flies as-such* but, rather, *sustenance* (food). Thus, perhaps the content is really something like *food* or *frog food*. If we consider the frog's *visual* mechanism — what it *sees*, i.e. — however, it is less clear that what it sees is *frog food*. For example, it seems that any number of appropriately sized, dark, appropriately moving, objects could trigger the frog's tongue-lashing behavior. Thus, perhaps the better content attribution is something like *small dark moving thing*. Similar considerations abound. Take the following non-exhaustive examples of content proposals for the fort's inner mental state: *fly*, *frog food, small dark moving thing, small dark moving food, fly stage, undetected fly part,* and others.[19]

The point of the frog-fly case is to illustrate the problem in providing a basis for the attribution of determinate contents of a mental representation among different compatible options that all account for the behavior of the system to which they belong (Smortchkova et al., 2020). The difficulty of providing such a basis is a problem which has faced a variety of theories, some of which I will outline below. Finally, DAMR recognizes the content determinacy constraint as a widely-accepted constraint, and purports to offer a solution to the

[17] Evaluating DAMR with regard to the myriad other formulations of this constraint is, furthermore, beyond the scope of this thesis.

[18] Stemming originally from work by Lettvin et al. (1959).

[19] See Neander (2006:168, footnote 3) for a useful, partial compendium of various content ascriptions — along with their classic defenders — related to the classic frog-fly case here described

problems faced which previous accounts have been unable to do (Egan, 2018, 2020). Whether or not DAMR does in fact underwrite the attribution of determinate content forms the central problem I address in this thesis, and I will argue in chapter 4 (4.2.) that there is reason to question whether or not DAMR is, in fact, able to answer to this constraint.

Attempts to Answer to the Content Determinacy Problem

There is no widely-accepted, agreed upon basis for the attribution of determinate representational content in an account of mental representation. However, numerous accounts have been proposed to answer this question (Shea, 2018; Egan, 2018, 2020). In this section I will focus on explicating several such accounts. The accounts I will explicate are so-called "tracking theories" of mental representational content, along with the problems they face. I t is important to note that the theories presented are *not* intended to be an exhaustive outline of available tracking theories or of all theories of mental representational content in cognitive science. Rather, I stick close to Frances Egan's (2018, 2020) own exegesis and focus on theories for their background relevance to DAMR, along with their broader influence.[20] This way of approaching such theories will serve to ensure that the theories discussed are relevant to Egan's motivations for DAMR, as she takes the failure of these theories to underwrite determinate content thus far as a central motivating factor for formulating DAMR.

Generally characterized, naturalistic tracking theories conceive mental representational content as a matter of a privileged relation between a given tokened internal mental state and some thing (typically, an object or property) which that state represents (Mendelovici, 2013; Egan, 2020). This privileged relation thus "tracks" the thing of which it is a representation. These theories are *naturalistic* in the sense that they conform to the (broadly) basic scientific assumptions about the world (they don't *presuppose* intentional capacities or representations).

One variety of tracking theory accounts are *information-theoretic* accounts (Dretske, 1981; Fodor, 1990). In broad terms, information-theoretic accounts hold that an internal state

[20] Another approach to providing a basis for determinate content which, like DAMR, also departs from seeking a privileged naturalistic relation is the "phenomenal intentionality" approach, which claims that what provides such a basis for determinate content has to do with a cognitive subject's *phenomenal experience* (see, for example, Horgan and Graham, 2012). Notably, Egan rejects such theories. It is beyond the scope of this thesis to provide a detailed reconstruction of her reasoning, but for her own discussion see Egan (2014:132-133).

S means *x* if and only if S is caused by the presence of an *x*, and certain further conditions — which aim to allow for the possibility of *mis*representation — obtain.[21] For example, to allow for the possibility that some S-tokenings are caused not by *x*s, but by something else *y* — in which case that S-tokening misrepresents a *y* as an *x*. Information-theoretic accounts face problems underwriting content determinacy, however, namely, in that everything in the causal chain from the distal object x to the internal S-tokening appears to satisfy the condition (Egan, 2020).

A second variety of tracking theory accounts are *teleological* theories (Millikan, 1984; and more recent versions in Ryder, 2004; Neander, 2017; Shea, 2018). In broad terms, teleological theories hold that internal state S means *x* if and only if S has the natural function of indicating cats. For example, on the teleological theory, some internal state means *fly* if and only if that state has the natural function of indicating flies. A central issue facing teleological theories again has to do with the content determinacy constraint: namely, that it is far from clear what natural function a given state possesses, and hence, that indeterminate functions cannot underwrite determinate content (Egan, 2018, 2020).

A third variety of tracking theory are *structural similarity* or *isomorphism/ homomorphism* theories (Cummins, 1989; Ramsey, 2007; Shagrir, 2012). Structural similarity theories generally hold that there is a relation akin to that relation which holds between a map and what it represents (i.e. a kind of *isomorphic* relation). There is a similarity, or "S-relation", between the elements of the map and the domain that is the target of the map (Gładziejewski and Miłkowski, 2017). A central problem faced for structural similarity theories is that, for any given set of internal states, there is likely to be many structurally-similar external conditions. Thus, there are also questions as to whether structural similarity theories can be constrained to only some, particular external conditions while still upholding naturalism (Egan, 2020).

Egan notes that tracking theories (namely, those addressed above) all face significant, outstanding problems when it comes to answering to the content determinant constraint

[21] Egan also identifies as a constraint that the account must provide for the possibility of *misrepresentation*. That is, for a given state to *represent*, it must be able to *mis*represent (Dretske, 1986). As Egan (2020) notes, the consequences of failing to meet the content determinacy constraint include failing to meet the misrepresentation constraint, since there is a *constitutive connection* between the two. I will not discuss this issue in detail in this thesis, but in chapter 5 I will suggest that this constitutive connection further qualifies the need for future research in relation to the worry I raise over content indeterminacy in chapter 4.

(Egan, 2020:7). This state of affairs will factor into a central theoretical motivation for DAMR. I will outline this motivation, along with DAMR's other central motivations, below.

DAMR's Theoretical Motivations

As I previously noted, tracking theories face outstanding problems when it comes to answering to the content determinant constraint. Importantly, Egan sees the continued employment of representational language in cognitive neuroscience *despite* the failure of tracking theories to provide the basis for determinate content as creating a kind of "mismatch". This apparent mismatch constitutes one of the central motivating factors for DAMR.

In addition to the apparent mismatch, DAMR also takes motivation from Noam Chomsky's (1995) eliminativism.[22] In broad terms, Chomsky's eliminativism holds that the representational states posited in accounts of cognitive capacities are not correctly construed as being about represented objects or properties, etc., but rather, are best construed as about sorting structures into kinds determined by the role they play in information-processing (Egan, 2014). On this view, characterizing a posited structure as 'representing' is just informal presentation which doesn't fit into the cognitive "theory proper" (Chomsky, 1995).

DAMR ultimately builds on two central ideas Chomsky presents: (1) that the computational theory proper excludes a representational characterization; and, (2) that talk of representing is a kind of informal presentation. However, as I will explicate in detail in chapter 3, DAMR departs from Chomsky's eliminativism in that it takes representational talk to be *uneliminable*. Instead, far from being subject to elimination, representational talk will, according to DAMR, play an essential role in computational cognitive neuroscientific explanations (this is why DAMR is *deflationary* rather than *eliminative*) (Egan, 2018, 2020).

DAMR is, thus, theoretically motivated by the combination of two central background factors. One motivating factor is the observation of the apparent "mismatch" between the continued employment of representational attribution in cognitive neuroscience and the simultaneous failure of classic tracking theories to provide a basis for such attribution. The

[22] William Ramsey (2020) has suggested Chomsky's view is *deflationary* as well, and refers to it as a kind of "non-relational" representationalism. Meanwhile, Egan has referred to Chomsky's eliminativist view as "Ersatz Represenationalism" (Egan, 2014, 2018).

other motivating factor is Chomsky's eliminativist view of representations posited in cognitive science. It is a motivating factor because it offers a promising alternative account of representational content to the naturalistic tracking theories. Ultimately, DAMR will take the central notions from Chomsky's eliminativism (as outlined above) and formulate an account which purports to provide a basis for the continued content ascription practices of computational cognitive neuroscientists in a way naturalistic tracking theories have been unable to.

3. Egan's Deflationary Account of Mental Representation (DAMR)

In this chapter I will explicate Frances Egan's (2020) deflationary account of mental representation (DAMR). First, in section 3.1., I will outline DAMR in articulation with what Egan refers to as the "computational theory proper." To this end, I will discuss Egan's construal of computation along with the five specifications involved in the theory proper. Second, in section 3.2., I will explicate DAMR's construal of representational content. To this end, I will explicate the two varieties of content DAMR maintains: mathematical content and pragmatic, cognitive content. Third, in section 3.3., I will explicate DAMR's realist construal of representational vehicles as *whatever* states or structures are specified and assigned content. Finally, in section 3.4., I will explicate how DAMR purports to meet the content determinacy constraint (outlined in chapter 2, section 2.2).

3.1. DAMR and The Computational Theory Proper

In this section, I will explicate DAMR with a primary focus on the "computational theory proper." First, I will articulate the computational theory proper in contrast to the pragmatic account of representational content (which will be explicated later, in section 3.2.). This will serve to specify just what the "theory proper" is supposed to be. Second, I will lay out the basic elements of physical system computation according to DAMR, which will serve to articulate the various components of DAMR in this section and the later sections of this chapter. Finally, I will explicate the computational theory proper in terms of its two main components: the *computational component* and the *ecological component*. I will first explicate the *computational component*, which consists in four computational specifications. I will then explicate the *ecological component*, which consists in an additional specification.

Broadly put, DAMR is an account of mental representation in computational cognitive neuroscience that couples a realist construal of representational vehicles with a pragmatic account of representational content.[23] [24] The *computational theory proper*, meanwhile, is the computational characterization of a cognitive system or device.[25] The pragmatic account of representational content says that mental content — what Egan calls cognitive content — is fundamentally *pragmatic*. That is, it is practice-dependent and serves a variety of pragmatic functions (detailed in section 3.2.). Cognitive content is ascribed by the computational cognitive theorist in an "intentional gloss" which is explanatorily useful but inessential to the full, scientific characterization of the success of the cognitive mechanism in question. Thus, the intentional gloss is a practice-dependent, glossy characterization of the actual, mathematical, "theory proper" characterization of the cognitive device or system in question.

---

[23] Egan has previously extended her pragmatic construal of representational content to, for example, computational cognitive psychology as well (Egan, 2009, 2014). The key point is that she excludes fields which she describes as "continuous with our commonsense conceptions of representation" (Egan, 2014). She does this because her aim is to provide an account of mental representation in cognitive neuroscience — *not* an account of representation *generally* (outside the specified scientific domain, i.e.). I consider the more restricted notion of DAMR, where it is centrally concerned with cognitive neuroscience, as this helps to constrain the discussion and because it is Egan's most recent articulation of DAMR (2020).

[24] In addition to the content determinacy constraint (outlined in chapter 2.2.), DAMR also wants to meet what Egan identifies as three other widely accepted constraints on an account of representational content (including the "misrepresentation" constraint outlined in footnote 20, chapter 2.2.). Furthermore, Egan also identifies William Ramsey's (2007) five "adequacy conditions" on a theory of mental representation as a useful set of conditions for evaluating DAMR overall (i.e. not just in terms of *content*). It is not the aim of this thesis to elaborate on these conditions or address how DAMR purports to answer to them. However, I will point to undertaking just such an evaluation as an important point for future research to address.

[25] Strictly speaking, as I will explicate in detail in this section, the *computational theory proper* also includes the "ecological component"; the computational characterization is the "*function-theoretic*" characterization.

I will explicate the intentional gloss and pragmatic construal of content in section 3.2. In this section, I will focus on the computational theory proper.

Before explicating these elements of DAMR in detail, it will be helpful to specify how Egan thinks a system computes a function in the first place. It will be helpful because the mappings, or *functions*, she identifies form the basis of how representational vehicles and contents are explicated (their relation). Egan identifies three basic elements for a physical system to perform a computation. A physical system computes some function *f* just in case there exists the following:

1. A mapping (*realization function,* f_R) from a system's physical state (or structure) types to vehicle types (i.e. a vehicle-type specification)
2. Physical state types (identified in f_R) related by a causal transition relation
3. A mapping (*interpretation function,* f_I) from a system's vehicle types (identified in f_R) to content (i.e. a content-specification identifying the values and arguments computed)

(1), f_R, specifies the representational *vehicles* of the system or mechanism computing the function. (2) specifies the physical, causal relations among the physical state types of the system. (3), f_I, specifies the representational *content* of the function computed. This is the basis of physical system computation on which DAMR aims to account for the nature and function of mental representation.

With this in mind, I will now turn to explicate the computational theory proper according to DAMR. The computational theory proper consists in the computational characterization of a cognitive system or device. It is the "theory proper" in that it suffices to completely explain the cognitive capacity in question.[26] Thus, in contrast with the intentional gloss, the theory proper is a practice-*independent*, mathematical-computational characterization of the cognitive device in question. It suffices to completely explain the cognitive capacity in question because it entirely explains the system's successes (and occasional failures) at a given cognitive task. What defines a given cognitive capacity will ultimately turn out to be a matter of our pre-theoretic explananda — this will be illuminated in

[26] In other words, it is the "theory proper" because it *comprises* the actual theory.

the section on content and the intentional gloss (3.2.). What is central to be clear about in terms of the theory proper is that the intentional gloss plays no role within the theory proper — it is, rather, merely a gloss *on* the theory proper.

The theory proper consists in five specifications (Egan, 2018). These five specifications can be further subdivided into two components: the computational component and the ecological component. The computational component consists of four of the five specifications. These four specifications are as follows:

1. The mathematical function computed by the mechanism
2. The algorithms involved in the computation of the mathematical function
3. The structures that the algorithm maintains
4. The computational processes defined over these structures

Together, these four specifications comprise the computational component of the theory proper. These provide the environment-independent characterization of the device and provide the basis for predicting and explaining the device's behavior in *any* environment. Egan also refers to the computational characterization of a cognitive device or system as *function-theoretic* (FT) (Egan, 2017, 2018). FT provides an *environment-neutral* and *domain-general* characterization of the device (Egan, 2014, 2017). It provides such a characterization because it *prescinds* any cognitive capacity in question and the environment in which the mechanism is typically embedded (Egan, 2018).[27]

The ecological component, meanwhile, is the fifth specification of the theory proper. The ecological component consists in general facts about the device's normal environment that explain why computing a specific mathematical function (whatever it is) suffices for the successful exercise of the cognitive capacity to be explained. The reason the ecological component is part of the theory proper is because a certain mathematical function will enable the organism to successfully exercise a given cognitive capacity only in a certain environment. Since the cognitive capacity (whatever it might be) is the theory's explanatory target, such environment-specific facts are required to provide an explanation of how the given computations involved contribute to the exercise of that capacity (Egan, 2018).

[27] The domain-specific or environment-non-neutral characterization is fulfilled by what Egan (2020) calls an "intentional gloss". The precise nature and function of the intentional gloss is detailed in section 3.2.

To better elucidate the ecological component, I will illustrate using an example borrowed from Egan: Marr (1982:337) identifies the relevant computational-mathematical characterization of a cognitive device's retinal image filtering process as computing a Laplacean convolved with a Gaussian (Egan, 2014:120). Consider the case of Computing the Laplacean of a Gaussian will only allow a given organism to detect an *edge* (i.e. to successfully exercise a certain visual cognitive capacity) in an environment with a *particular* physics. In a physically different environment, computing the Laplacean of a Gaussian might not help the device to *see* at all. The ecological component is also important for Egan's account of representational *cognitive* content (Egan, 2018). I will return to this point in section 3.2., where I explicate cognitive content.

Taken together, the computational component and the ecological component — consisting in the five specifications outlined — comprise the computational theory proper. According to DAMR, this is entirely sufficient to explain the cognitive mechanism.

3.2. Representational Content and the Intentional Gloss

Having explicated the computational theory proper in 3.1., I will now turn to explicate representational *content* on DAMR. I will do this by distinguishing DAMR's two kinds of content: mathematical content and cognitive content. I will first explicate mathematical content, paying particular attention to why it is essential to the theory proper. Then, I will explicate cognitive content. I will do this first by showing how cognitive content is characterized in rejecting two commitments tracking theories hold. Second, I will show how, in rejecting those commitments, cognitive content turns out to be inessential to the theory proper, but essential as an "intentional gloss" on the theory proper fulfilling broadly *explanatory* or *pragmatic* purposes. Finally, I will outline the various pragmatic functions cognitive content maintains.

Before I explicate representational content on DAMR, it will be important to recall Egan's motivation for DAMR. This will serve to partially justify the account sketched. One of the central motivating factors for DAMR is making sense of the theoretical mismatch (see 2.2.): the continued ascription of representational content to posited states and structures in cognitive neuroscientific practice *despite* the fact that naturalistic tracking theories have failed

to meet certain widely accepted constraints on content. Egan maintains that cognitive neuroscientists must use such ascriptions for *some* reason — that their practice is not "*wrong*" or *eliminable*, in other words. Egan sees this combination of factors as justifying the alternative approach, DAMR, to be exposited. With that in mind, and having distinguished between the theory proper and DAMR in 3.1., I will now turn to explicate representational content on DAMR.

What Egan calls an *interpretation function* — f_I — maps content from the representational vehicles specified via f_R. To this end, DAMR distinguishes between two distinct kinds of representational content: *mathematical* content and *cognitive* content. I will explicate each in turn.

Mathematical content is representational content which represents the arguments and values computed by a cognitive system. For example, Marr's (1982) theory of early vision purports to explain edge detection by positing computation of the Laplacean of a Gaussian. Intensity values are taken as inputs by the mechanism at points in the image, and the mechanism then computes the rate of intensity change over the image and outputs the value of a mathematical function (here, the Laplacean of a Gaussian) (Egan, 2018).[28] Such inputs represent *arguments* and such outputs represent *values* (of the mathematical function of the cognitive task in question, i.e.). Thus, the computational characterization specifies a kind of content, namely, *mathematical* content. DAMR holds that mathematical content is *essential* to the theory proper: *if* the mechanism computed a different function, *then* it would be assigned different mathematical contents (hence, it would be a different mathematical state) (Egan, 2018).

I will now turn to characterize the other kind of representational content distinguished by DAMR: *cognitive* content. Egan (2020:11) characterizes cognitive content as the sort of content philosophers "typically have in mind" when discussing representational content. To do so, it will first be helpful to note that DAMR builds on a central feature of tracking theories. Recall that tracking theories hold that states of mind represent aspects of the world by *tracking* the objects (properties, etc.) they are about.[29] DAMR, however, rejects two central

[28] Other examples include: addition, vector subtraction, and a fast Fourier transform, e.g. (Egan, 2018).

[29] For my detailed explication of tracking theories, and in particular, the tracking theories which form the theoretical backdrop of Egan's formulation of DAMR, see the chapter 2 (2.2.) subsection on naturalistic attempts to answer the content determinacy constraint.

commitments tracking theories shares. One of these commitments (1) is the commitment to mental representations having their contents *essentially*. That is, that a given mental representation has its content such that, if it had some other content, it would be a different mental representation. The other of these commitments (2) is the commitment to meeting the content determinacy constraint via a privileged naturalistic relation holding between a state or structure and the object or property it is about (Egan, 2018). As I outlined in chapter 2 (2.2.), this forms the basis for the various tracking theories' attempts to meet the content determinacy constraint.

DAMR rejects (1) because a given mechanism characterized computationally on the theory proper could subserve a different cognitive capacity if it were embedded in the system differently. Its representational vehicles, then, would be ascribed a different cognitive content. Thus, cognitive content turns out to be *inessential*. For example, suppose a given computationally characterized mechanism subserves a *visual* capacity. In such a case, the mechanism will be ascribed *visual* cognitive content in f_I (since the cognitive capacity in question is a *visual* capacity). But suppose the mechanism were embodied in the system differently, such that it no longer subserved a visual capacity but, rather, an auditory one. Then, now serving an *auditory* capacity, the mechanism would be ascribed *auditory* cognitive content. Meanwhile, DAMR rejects (2) because it holds that content determinacy is secured via the explanatory context along with certain other pragmatic considerations (Egan, 2014, 2018, 2020). I detail these considerations below.

The upshot of DAMR's rejection of these commitments is that a given state or structure does *not* have representational content *independent of* the interpretive practice of theorists, to whom it serves various pragmatic aims. For these reasons, Egan construes content as a kind of *gloss* on the computational theory proper, namely, an "intentional gloss". Such content attributions are *intentional* in that they employ intentional language and descriptions. They constitute a *gloss* in that they serves purely pragmatic roles in service of the explanatory concerns of a computational cognitive theorist and are inessential to the theory proper.[30] The

[30] As a consequence of cognitive contents inessentiality with regard to its representational vehicles, Egan (2014) notes that such vehicles (characterized in computational cognitive science) do *not* seem to have so-called "original" or "intrinsic" intentionality. This is in contrast with the commonsense view of mental states (which are traditionally construed as having *intrinsic* as opposed to non-intrinsic, "derived" intentionality). It is beyond the scope of this thesis to draw out the consequences of this conclusion in detail. For clearer formulations on the original/derived intentionality distinction see Searle, J. (1980, 1993) and Egan's (2014) conclusions with regard to DAMR.

intentional gloss involves, hence, *domain-specific* content. In other words, it is content that is specific to the cognitive capacity to be explained (Egan 2010, 2020).

With relation to the theory proper, the gloss serves as a kind of helpful explicator. For example, the structures that a given algorithm maintains in the theory proper will be *glossed* as representations, and the computational processes defined over the structures that the algorithm maintains are glossed as representational processes. In broad terms: the computational theory proper is ascribed cognitive content in an intentional gloss (in f_I). In this way, cognitive content is inessential to the theory proper, as described. However, cognitive content and the intentional gloss are essential for another reason. The reason is this: they serve to fill or bridge the pre-theoretic *explananda* with the theoretic (theory proper) *explanantia*.[31] Thus, unlike representation in Chomsky's eliminativism which holds that representational talk is ultimately subject to elimination from cognitive scientific practice, DAMR holds that content is *not* eliminable from such practice — this is the sense in which DAMR is *deflationary* rather than eliminative.[32]

DAMR maintains at least three pragmatic functions which the cognitive content serves. These are as follows: First, as I have already highlighted, the intentional gloss links a theory's *explanantia* — the mathematical capacities posited in a computational theory, i.e. — with the cognitive capacity that is the theory's *explananda*.[33] In illustration, take Egan's own example. The ecological component of the theory proper might specify facts about covariation between tokenings of the structure in question and the distal property instantiations, and while these help explain an organisms's visual capacity they do not say anything about representation. The intentional gloss, then, is ascribed such that it shows how the theory addresses its explanatory target. It does this by characterizing the organism-environment

---

[31] Egan articulates this role in terms of answering to Wilfred Sellars's (1962) notion of the "explanatory gap" between scientific theories and our pre-theoretic, commonsense understandings of the phenomenon in question. Egan understands cognitive neuroscience and the theory proper to provide a *reductive* theory of pre-theoretic *explananda*. Thus, a la Sellars, there is a "gap" in explanation leftover between the two. Egan maintains it is the intentional gloss which is *essential* for filling or bridging this gap. For Egan's discussion see Egan (2020:26-27).

[32] DAMR is also distinct from two additional "anti-realist" (or *irrealist*) positions: *fictionalism* and *interpretivism*. Broadly put, fictionalism about mental representation holds that intentional talk genuinely aims to describe the world but doesn't suppose there are real representational objects (Sprevak, 2013). DAMR is not a kind of fictionalism, as Egan notes, because fictional objects *cannot* be causally efficacious (Egan, 2020:22). Interpretivism, meanwhile, differs from DAMR as typically construed (in Dennett (1987), per Egan's example).

[33] Egan sometimes also identifies mathematical capacities posited in the computational as "subpersonal level" capacities and the pre-theoretic explananda as "personal level" capacities. This personal/subpersonal level distinction is not the subject of evaluation in this thesis, but for further discussion about the distinction see Drayson (2014).

interaction that enables a given cognitive capacity in terms of representation, (Egan, 2014, 2018, 2020).

Second, the intentional gloss helps keep track of the “flow of information" in a given cognitive system. “Flow of information” here means changes in a system caused by internal processes and external (environmental) events. So, the intentional gloss helps theorists follow the changes undergone in a given cognitive system where the purely computational characterization might be difficult to follow. Egan notes that the choice content ascription will be sensitive to things like *ease of explanation*. For example, an idealized, conspicuous explanation will serve to help *theorists* more readily track a systems’s changes than an opaque, purely computational description will.

Third, content ascriptions can serve as a placeholder for an incompletely developed computational theory. In this way, the intentional gloss can serve to guide the further development of the theory (Egan, 2020). Importantly, Egan clarifies that content ascriptions which serve this placeholder function are not appropriately construed as a gloss but, rather, as playing a “gloss-like” role. I will return to consider this pragmatic function of content in chapter 4 (4.2.) where it is central to the theoretical concern I raise about how DAMR meets the content determinacy constraint (outlined in 2.2.).

3.3. Representational Vehicles

In chapter 2 (2.2.), I outlined the distinction between representational content and representational vehicles. In the previous section (3.2.) of this chapter, I explicated DAMR’s account of representational content, including the *pragmatic* account of cognitive content. In this section, I will explicate DAMR’s accompanying *realist* construal of representational vehicles. Before before doing so, I will first recall what representational vehicles are supposed to be, then, I will add an important qualification necessary for properly understanding the representational vehicles DAMR maintains.

Recall that representational vehicles are the carriers or bearers of representational content. They are *physically realized* states or structures that carry or bear content. Since representational vehicles are physically realized, they are also *causally efficacious*. The kind of states or structures they are construed as depends on the specific theory in which they are

invoked. Traditionally, on CRTM, vehicles are *symbols* which carry representations. However, other construals — such connectionist network node activation patterns — have also been posited as vehicles.[34]

Having recalled what representational vehicles are supposed to be, there is an important qualification I will now make regarding DAMR's construal of representational vehicles. It is this: DAMR explicitly seeks to provide, in its account of representational vehicles, an answer to the question of what it is to function as a representation *in an explanatory account of a cognitive capacity*. In other words, it seeks to provide an answer to the question of what it is to function as a representation in an account of a cognitive capacity, *not* — as some approaches to representation do — as a representation *in general*. Thus, it would be, for example, inaccurate to construe DAMR as claiming to provide an answer to the question of what representations are in a generalized, metaphysical sense or in a sense that encompasses all our commonsense uses of the notion of representation (Egan, 2020).

With this qualification in mind, I will now turn to explicate DAMR's construal of representational vehicles. According to DAMR, representational vehicles are states or structures identified in the theory proper via f_R and assigned contents via f_I. Thus, representational vehicles are — as opposed to cognitive content — as *real* as any states or structures posited in scientific theories. This is the sense in which representational vehicles are construed along *realist* lines in DAMR. Whereas cognitive content is an ascription made in an intentional gloss, representational vehicles *are* the states or structures to which cognitive content is so ascribed. Representational vehicles are not something "over and above" the posited states or structures, rather, they *are* such states or structures grouped in a certain way such that they are identified via f_R and assigned content in f_I.[35]

To better understand what constitutes vehicles on DAMR, it will be informative to use Egan's (2020) own example of *genes* in molecular biology. Genes are physically realized in chemical structures. These structures (analogous to the physical states or structures identified in cognitive theories) are grouped together via their causal efficacy in producing the proteins responsible for various given phenotypical traits. In this sense, they are "abstracted away"

---

[34] See chapter 2.2. for my more thorough explication of the vehicle-content distinction.

[35] Egan (2020) maintains that this means DAMR avoids Ramsey's (2007) condition that vehicles must be more than "mere causal relays." According to Egan, unlike mere causal relays which have no content, DAMR's vehicles are ascribed content in the intentional gloss.

from more basic, underlying physical properties. f_R similarly abstracts away from the more basic properties of the states or structures realized and groups them by their role in a given cognitive capacity. Thus, Egan contends, just as genes are construed along *realist* lines, so too should be the representational vehicles in construed in DAMR.

Finally, an important point regarding the nature of representational vehicles in DAMR is that it does *not* matter what *kind* of posited states and structures serve as representational vehicles. In other words, what states and structures can count as representational vehicles is quite open. For example, *any* cognitive model — whether classically symbolic, connectionist, etc. — can group together structures via f_R and assign them contents in f_I. Egan notes one exception: the relevant objects must be capable of having causal powers, since f_R specifies the causal organization of the system. Thus *abstracta*, or abstract objects, cannot serve as representational vehicles, since they do not have causal powers.

3.4. Answering to the Content Determinacy Constraint

In this section I will explicate how DAMR purports to meet the content determinacy constraint (outlined in 2.2.) which forms the basis of this thesis's research problem.[36] The content determinacy constraint requires an account provide the basis for the attribution of determinate content to posited states or structures. Egan variously suggests that the combination of the *explanatory context* (the cognitive capacity to-be-explained) and additional *pragmatic factors* (as explicated in 3.2.) successfully answer to this constraint. However, what precise variety of factors Egan takes as *sufficient* to provide the basis for the attribution of determinate content is somewhat unclear.

Broadly, there are at least three elements which taken together in some combination are supposed to provide the basis for the attribution of determinate content. First, there is what Egan (2014) has referred to as the "explanatory context". The explanatory context is the pre-theoretic explanandum of the theory in question. In other words, it is the cognitive capacity which the theory is supposed to explain. So, prior to the development of a given cognitive theory, there will be some question about how an organism *does what it does* — i.e., has some cognitive capacity which enables it to interact with its environment in whatever particular

[36] Egan is unclear about what counts as an "attribution of determinate content".

way. Thus, the explanandum — the cognitive capacity in question — is pre-theoretic because it precedes the theory to-be-developed. This element of the basis for attribution combination thus serves to focus what the theory will be about in broad terms.

Second, there are the various naturalistic relations. These tracking relations each characterize a way the cognitive state can fit the world. In the tracking theories I discussed in chapter 2, finding *the*, privileged, naturalistic relation is what it would take to answer to the determinate content constraint. However, as an element of content determination on DAMR, the naturalistic relations are all equal candidates which the theorist will select among or choose between given some motivating pragmatic considerations. Both the pragmatic considerations and the naturalistic relations will be constrained by the initial explanatory context.

Third, there are the additional, pragmatic considerations. This is a point DAMR is somewhat unclear: it is unclear *which* pragmatic considerations, *in what combination*, are supposed to secure determinate content attribution.[37] For example, Egan does *not* specify if any one of the pragmatic considerations she outlines is sufficient *on its own* to secure determinate content attributions, or if some particular combination is so sufficient, or if *all of the pragmatic functions together* are sufficient to determine content, e.g. However, she is explicit that *some variety of them* — along with the other elements discussed — *are* sufficient. She says, for example, "Once the role of specific explanatory interests and other pragmatic factors in content attribution is fully appreciated, determinacy is to be expected," (Egan, 2018:255). Hence, though it is unclear exactly *which* such considerations (along with the other elements mentioned) in what combination are sufficient to secure content determinacy (and hence meet the content determinacy constraint), Egan is clear that *some* such variety of considerations *are* sufficient.

---

[37] This points to a larger problem of unclarity or under-specification in DAMR, namely, how the various pragmatic functions *interrelate*. For example, Egan says some content ascriptions fulfilling some pragmatic functions as construable as an intentional gloss. Meanwhile, content ascriptions fulfilling other pragmatic functions (the "placeholder" function, e.g.) are merely "gloss-like." As a consequence, it seems that at least *some* pragmatic functions are incompatible with each other. They are incompatible in that a single content ascription could *not* serve both of them at the same time. To illustrate this, consider that a content ascription of some vehicle(s) fully specified in f_R (i.e., in a complete theory) fulfilling some pragmatic function *cannot* serve a *placeholder* function. This is because the placeholder function is defined as content ascribed in an incompletely developed theory where the representational vehicles are yet-to-be-specified. However, it is unclear if the reverse holds: *if* a content ascription serving as a placeholder *can* or *does* also fulfill other pragmatic functions (and, if so, which ones). It is beyond the scope of this thesis to fully explore how such interrelations might be fully specified. However, this is one area I will highlight in chapter 5 as warranting further research.

Now that I have highlighted a point of unclarity in DAMR's account of content determinacy, I will briefly recall the pragmatic considerations with Egan exposits and then outline how DAMR purports to meet the content determinacy constraint.[38] Pragmatic considerations which are highlighted by Egan include the following: salience or tractability, ease of explanation, filling the "explanatory gap", and acting as a placeholder in incompletely developed theories. The salience or tractability consideration means that the determination of content is sensitive to how *we* (*theorists*, i.e.) keep track of internal changes in the cognitive system in question. The ease of explanation consideration has to do with what makes sense for our theoretic practice. Filling the "explanatory gap" is the consideration which holds that content needs to explain the relationship between the reductive cognitive theory with which we ended up and the pre-theoretic explanandum with which we began. Finally, a content ascription can act as a "placeholder" in an incomplete theory and so guide its development.

With this in mind, the way that determinate content attribution is secured according to DAMR is via a specified version of something like the following: first, there is an explanatory context: that is, whatever the pre-theoretic explanandum happens to be. In cognitive neuroscience, it will be some cognitive capacity, e.g. Second, there is some variety of naturalistic relations which the explanatory context constrains. The explanatory context constraints such naturalistic alternatives because many naturalistic relations will end up outside the scope of the explanatory context. If the explanatory context is a frog's fly-catching mechanism (the cognitive capacity to-be-explained), then the cognitive theorist can already rule out some *other,* non-fly-catching mechanism and its possible associated content. In the case of a frog's fly-catching mechanism, there will be various naturalistic candidates like *fly*, *frog food*, *small dark moving thing*, etc. Third, there are the additional pragmatic considerations. These include the various pragmatic considerations previously outlined: salience, ease of explanation, and acting as a placeholder. Some variety of these pragmatic considerations will motivate the theorist to select from among the naturalistic alternatives and ascribe some representational content. Thus, according to DAMR, securing the attribution of determinate content and thereby allowing the account to meet the content determinacy constraint.

[38] For the full explication of the various pragmatic functions of content DAMR maintains, see 3.2.

4. Content Indeterminacy in DAMR

In this chapter, I will raise a theoretical concern for DAMR, namely, that one of the pragmatic functions of content it maintains might be vulnerable to content *indeterminacy*. Ultimately, in chapter 5, I will connect the conclusions of this consideration back to the central research problem outlined in chapter 1. To this end, first, in section 4.1., I will briefly outline several previous criticisms of DAMR and highlight the fact that typically criticisms have *not* focused on content determinacy. This will serve to contrast with the theoretical concern I will raise in section 4.2. In section 4.2., I will raise a theoretical concern for DAMR, namely, a concern which states that one of the pragmatic functions of content DAMR maintains might be vulnerable to content indeterminacy. I will also briefly illustrate how this might open up the broader concern that such indeterminacy could be generalized or extended to other pragmatic functions which DAMR maintains. Finally, in section 4.3., I will outline several possible responses to the concern raised in 4.2., and evaluate each in turn.

4.1. Previous Criticisms of DAMR

In this section I will briefly outline some previous criticisms of DAMR and articulate how these criticisms differ from the theoretical concern I will raise in 4.2. DAMR has previously been criticized for a broad variety of reasons. I will not explicate the criticisms to follow in detail, rather, I outline some of the ways in which DAMR has been previously criticized so as to articulate how the concern raised in the following section (4.2.) differs from those previously raised.

Some criticisms and theoretical concerns regarding deflationism, such as Bechtel (2016), have focused on cognitive scientific *practice*, arguing that deflationary accounts fails to successfully describe the role representation plays in cognitive scientific practice (in Bechtel's case, that cognitive content is, in fact, *essential*).[39] Coelho Mollo (2017) assesses a line of criticism against DAMR which maintain that DAMR ultimately collapses into a more traditional view, such as eliminativism, e.g. (that DAMR does not provide adequate reasons to

[39] Egan has responded directly to this criticism — see Egan (2020).

think cognitive content really is *essential*)[40] More recently, Hutto and Myin (2020) have argued that mathematical content makes *no causal difference* in a cognitive system. To this end, they further argue that mathematical characterizations of such a system are thus, at best, a *mathematical gloss* on that system's behavior — having similar pragmatic and heuristic values as Egan supposes the *intentional gloss* to (i.e., that mathematical content, too, is *inessential* to the theory proper). Meanwhile, Ramsey (2020) has argued, in response to DAMR, that cognitive content *is essential* in the exercise of a cognitive capacity *even if* embedding the mechanism differently would produce different content.

As I have just illustrated, DAMR has faced criticism in a variety of ways. However, most criticisms have focused on the *essentiality* of representational content (for cognitive content: Bechtel, 2016; Ramsey, 2020; the criticisms outlined by Coelho Mollo, 2017; and for mathematical content — Hutto and Myin, 2020). That is, whether or not a given variety of representational content identified in DAMR is *essential*. The theoretical concern I raise in 4.2. will differ from these previous concerns and criticisms in that it will focus — *not* on the essentiality of cognitive or mathematical content — but rather, on the *content determinacy constraint* and, thus, how DAMR meets this constraint.

4.2. A Content Indeterminacy Concern for DAMR

In this section, I will raise a theoretical concern for DAMR. Namely, I will raise a concern related to how DAMR purports to meet the content determinacy constraint (as outlined in 2.2.). Ultimately, in chapter 5, I will connect the conclusions I draw from raising this concern back to the central problem of this thesis: assessing if DAMR meets the *content determinacy constraint* on an account of representational content. First, I will reiterate what the content determinacy constraint is, and emphasize in what way I understand the constraint (in particular, with regard to the two possible understandings of the constraint outlined earlier, in 2.2.). Second, I will introduce an analogy which will ultimately serve to help motivate the theoretical concern raised. Third, I will raise a concern which states that one of the pragmatic functions DAMR maintains (namely, the "placeholder" function (PF)) might be vulnerable to content indeterminacy. Fourth, I will illustrate with a thought experiment how such content

[40] Egan (2020) is also sure to differentiate DAMR from eliminativist views, as I have explicated herein (see chapter 3.2., e.g.).

indeterminacy might arise for PF. Finally, I will conclude by outlining a broader concern, namely, that the concern raised for PF might be *generalizable* or *extended* to other pragmatic functions DAMR maintains.

I will now reiterate what the *content determinacy constraint* is. The content determinacy constraint states: "The account should provide the basis for attributing *determinate contents* to computational states or structures," (Egan, 2020). Centrally, DAMR is committed to meeting this constraint (Egan, 2018, 2020). As outlined in chapter 2 (2.2.), I will assume this constraint is about the attribution of a certain kind of content, namely, *determinate* content. As a consequence, the concern I raise should not be taken to apply to the other, second understanding of the content determinacy constraint outlined in chapter 2 (2.2.), on which the constraint is understood as being primarily about *attribution*. However, I will address the second understanding in section 4.3. when I evaluate the concern I here raise; and, in chapter 5, where I will point to future research directions.

Before I raise a theoretical concern for DAMR, consider the following everyday case of fulfilling a pragmatic function. Imagine the following case: I need to paint my kitchen table. I purchased it unpainted several weeks ago, and it's finally time to stop procrastinating and get the job done. In order to accomplish this task, there are a variety of things I need to use: I need to use *paint*, for example, and I need to use a brush, and I need a bucket, other such usual components, etc. Importantly, these things serve *pragmatic* functions for me: a brush lets me apply paint, a bucket lets me have easy access to the paint, and paint allows me *to* paint. However, notice that — in serving those pragmatic functions — it doesn't matter which *brand* of paint I use. Nor does it matter which *color* paint I use. Nor does it matter the size of the bucket, so long as it's both big enough and not too big: there's many different sizes the bucket could serve the pragmatic function for which I need it. The take away is this: many *different* things could equally well serve to accomplish the *same* pragmatic function — the pragmatic function of the paintbrush could be equally well met by various other different kinds of paintbrushes, e.g.

I will now raise a theoretical concern for DAMR's answer to the content determinacy constraint which is analogous to the case I just described. The concern is this: at least one of the pragmatic functions Egan outlines might be vulnerable to content *indeterminacy*. By this I mean that the function might *not* provide the basis for the attribution of determinate content

(one content rather than another, i.e.). Before I detail this concern, there are two important clarifications I want to make about what's relevant in the painting analogy above and the case of representational content which I will address. The first clarification is this: it is *not* relevant, for my purposes, what exactly constitutes or defines "pragmatic". Even if the case I described, or the pragmatic function Egan highlights, turned out not to be *pragmatic* (in whatever sense), the relevant issue I will raise would still stand: multiple different things can sometimes accomplish one and the same task or goal. The second clarification is this: I will *not* focus on every one of the pragmatic considerations Egan highlights. Consequently, I will *not* conclude that the concern I raise applies to all of them.[41] Rather, I will limit my focus to the particular pragmatic function of representational content ascription already stated, namely, PF. Hence, I will consider PF in isolation from the other pragmatic functions Egan outlines.[42] This will serve to focus my discussion.

First, it will be helpful to get clear about what Egan means what she says that one of the pragmatic functions of content is to serve as a "placeholder" in an incomplete theory. Egan states the following:

> "A content ascription can serve as a temporary placeholder for an incompletely developed computational theory of a cognitive capacity and so guide the discovery of mechanisms underlying the capacity. For example, at the early stages of theory development, prior to the specification of the mathematical function computed and the structures and processes that enable the computation, a visual theorist may characterize a to-be-specified structure as representing edges or some other visible property of the distal scene," (Egan, 2020:13).

> "Characterizing to-be-discovered structures in terms of content allows the theorist to formulate hypotheses about the causal roles of the structures she is investigating," (Egan, 2020:30).

---

[41] I *will* suggest, in chapter 5, that *if* and *to what extent* similar concerns could be extended to other pragmatic functions of content is a question which could form the basis for further clarification in future research, however.

[42] Whether or not PF is isolated from all the other pragmatic functions according to DAMR is unclear. As I stated previously in 3.3., I will raise this point as an area for further research to address. However, regardless of whether or not it is always isolated so on DAMR, limiting my discussion to one pragmatic function will still serve to helpfully focus the concern raised: if a particular pragmatic function admits content indeterminacy, i.e.

So, for a content ascription to serve as a "placeholder" prior to the specification of the theory proper, it needs to: serve to help a theorist formulate hypotheses about the causal roles of the structures she is investigating, and in this way serve to guide the theorist's development of the computational theory proper.

I will now detail the theoretical concern about content ascription which serves the pragmatic function of being a placeholder in an incomplete theory. The concern is this: just as I can paint a table equally as well with different kinds of paintbrushes, so too it seems that there could be multiple different content ascriptions which could equally well serve PF, and this might make PF vulnerable to content indeterminacy.[43] [44] To see how this could come about, first consider what PF needs to accomplish, per Egan: it needs to allow the theorist to formulate hypotheses about the causal roles of the structures she is investigating (and in this way, so guide her research). However, it seems that more than one content ascription could *equally well* allow the theorist to formulate such hypotheses. I will elaborate this idea below.

To draw on the classic fly-frog example (which Egan herself also returns to): *small dark moving thing* is a *visual* content ascription for a frog's tongue-lashing mechanism (the cognitive capacity to-be-explained). But, so far as PF is concerned, *small dark moving food* could equally well serve as a content ascription. It, too, is a visual content ascription and could equally well allow a theorist to formulate hypotheses about the mechanism in question. In this way, either content ascription — *small dark moving thing* or *small dark moving food* — could fulfill PF just as well: both are just as suitable for guiding hypotheses about the underlying mechanisms involved.

With this consideration in mind, consider the following scenario: suppose there are two cognitive theorists, Theorist A and Theorist B. Further suppose each of them is working

---

[43] What exactly constitutes a "difference" between two or more content ascriptions is beyond the scope of this thesis to fully investigate. Shea (2018:82), for example, points to a "difference in content" consisting in contents which "could come apart." To what extent this question has been clarified is a fruitful area for further research. However, in what way such content could come apart is left unspecified. My strategy will be to follow Egan's own examples closely to ensure that, whatever a difference-in-content amounts to, the examples used here are as genuinely *different* (in *whatever* sense, i.e.) as those employed in Egan's explication of DAMR.

[44] The painting analogy I laid out is supposed to motivate the more general idea that multiple things can equally well fulfill a given pragmatic function. This generality is meant to make conspicuous how such a concern might be generalized or extended to other pragmatic functions (in DAMR, i.e.). However, the analogy could also be articulated in a "placeholder-specific" way: suppose I want to paint the table blue, but I don't yet have access to blue paint. I can already start by giving it one under-coat with *whatever* color I do have — it will do to serve as a placeholder until I get the proper paint, e.g.

on specifying the underlying structures involved in a frog's tongue-lashing behavior. Let's also suppose — to ensure that their content ascriptions are "placeholder" ascriptions — that they are working at a time before anything much is known about how this behavior works (PF is, thus, *gloss-like*). In order to better keep track of mechanism, A might ascribe *small dark moving thing* (call it "$content_a$") while Theorist B ascribes *small dark moving food* (call it "$content_b$") as representational content, in specific, to yet-to-be-specified structures which underlie the mechanism in question.

I want to suggest this take away from the above scenario is this: both Theorist A's and Theorist B's content ascriptions equally well fulfill the same pragmatic function in question, namely, PF. This is because both content ascriptions enable the theorists to formulate hypotheses about the causal roles of the yet to-be-discovered structures, and this is just what PF functions to do. However, both Theorist A and Theorist B ascribe *different* content ($content_a$ and $content_b$, respectively) despite all other considerations being the same. Put another way, Theorist A and Theorist B both rely on the fact that different content ascriptions could *equally well* fulfill PF.

The apparent upshot of this take away is as follows: *different* representational contents both fulfilling the same pragmatic function (PF) are attributed to the same underlying states. Thus, it seems that *nothing in PF* determines $content_a$ rather than $content_b$. Thus, it appears as if the content of PF, in the case described, is *indeterminate*. It so appears this way because PF does *not* seem to provide the basis for the attribution of one content over the other. Consequently: PF appears vulnerable to content indeterminacy.

There is an important clarification I want to make regarding the preceding considerations. The first clarification is this: I am not suggesting that *there could be no* pragmatic consideration which *would* privilege one content ascription over another in such a case. Indeed, there could be yet another theorist whose own explanatory concerns involve or require making just such a distinction — placing an emphasis on downstream nutritional effects, e.g., and thus preferring the content *small dark moving food*. Instead, what matters here is that there are *some* cases — such as the one described — in which *no* such distinction is involved. It matters because it suggests that two theorists with the the same pragmatic consideration (PF) might, nonetheless, ascribe different contents, leading to the indeterminacy described.

Finally, I will briefly highlight a broader theoretical concern that might be implicated by these considerations. The concern is this: the content indeterminacy concern for PF might be *generalizable* or *extended* to other pragmatic functions in DAMR. If it were, then DAMR would face a broader challenge in meeting the content determinacy constraint. This issue is beyond the scope of this thesis to address. However, I will suggest that further clarificatory research is warranted.

4.3. Responding to the Indeterminacy Concern

In section 4.2. I raised a theoretical concern which stated that one of the pragmatic functions DAMR maintains might be vulnerable to content indeterminacy. In this section, I will outline five potential responses to the concern raised in 4.2., and evaluate each in turn.

One potential response to the concern raised in 4.2. is this: two theorists *cannot* ascribe different content to the *same* states (or structures) *if* they have *not* yet specified what those structures are. In other words, *if* a theorist has *not* yet posited a state to which they could assign content, then *there is no* specified state to which another theorist could assign different content. However, Egan seems to accept that the "placeholder" pragmatic function for content ascription in incompletely developed theories is a variety of content ascription which does not necessarily assign content in f_I. Thus, it might be true that at such an early stage in the development of a theory, the theorist has not yet specified the structures which will serve as representational vehicles in the intentional gloss. However, nonetheless, the theorist assumes there is *some* such structure there — *even if* it is yet to be fully specified. Thus, two theorists could presumably ascribe two different content (as in the case described in 4.2.) to the same states *whatever they turn out to be*.[45]

Another potential response is this: since such placeholder content ascriptions are *not* always part of the intentional gloss, it is unclear that any content indeterminacy in PF — even if it is so vulnerable — ought to be a concern for DAMR. In other words, if PF is "outside of" the intentional gloss, perhaps *it does not matter* if PF is vulnerable to indeterminate content.

[45] Compare, for contrast, that a similar "placeholder" ascription would *not* admit this variety of indeterminacy in a theory which construed *realist* representational content (which would typically assume a privileged naturalistic content-determining relation). This is because such a theory would assume there *is* some "correct" determinate content — so even if two theorists ascribe different content, it would be assumed that ultimately *at most* only one of them is correct.

However, Egan (2020) highlights PF as one of the pragmatic functions of content ascription: such content plays a role in cognitive neuroscientific practice *even if* it is not part of the intentional gloss. Thus, *if* the content indeterminacy concern holds for PF, *then* it is a concern for DAMR nonetheless.

Another potential response is this: perhaps there is, in fact, a content ascription which does "best" fulfill PF in every case — a *privileged* content ascription, i.e. Such an ascription might be, for example, *whatever* ascription will allow the theorist to formulate the best hypotheses, make the best predictions, or best complete the theory. However, this is precisely what is in question in the concern raised in section 4.2.: that there *is* a single content ascription most suited to such a task in the first place. Thus, this response leaves open the need to demonstrate that there will be, in fact, a single "best" PF content ascription.[46]

Another potential response is this: perhaps some degree of indeterminate content *is acceptable*. In other words, this response would say that the concern is successful, but that it is, nonetheless, *not* a problem for DAMR if it were to admit some degree of indeterminate content. Indeed, some philosophers find it unproblematic — and even *expected* — that accounts of mental representational content will admit some degree of content indeterminacy (Shea, 2018; Ramsey, 2020). However, a potential problem with this response is that, were DAMR to admit some degree of content indeterminacy, then it would not answer to one of its central motivations: the failure of tracking theories to provide a basis for the attribution of determinate content. As Hutto and Myin (2020) put it, this would be to fail to show that DAMR is "…a net explanatory gain." Even so, perhaps there would be reasons, nonetheless, to think DAMR is preferable to its competitors. It might turn out that, for example, DAMR admits *less* indeterminacy than its competitors, or admits indeterminacy in some less arbitrary, principled (i.e.) way.

Finally, another potential response is this: suppose that we conclude (from the considerations in 4.2.) *not* that content is *indeterminate*, but rather, that the better understanding of DAMR's commitment to the content determinacy constraint is the other, second understanding outlined in chapter 2, section 2.2. — namely: that the content determinacy constraint is foremost about *attribution* (about something theorists call "determinate content"), and that it does *not* matter if the content attributed is, in fact,

[46] It would need to show that there would be such an ascription *in every case* — if there were *not*, then there would be at least *one* case of indeterminate content a la the variety described in 4.2.

determinate or not. In other words, perhaps the best response to the concern is to investigate if DAMR needs to be committed to the understanding of the content determinacy constraint assumed. As previously noted, I contend that this line of response warrants further investigation. As I have outlined (in 2.2.), it is unclear how exactly DAMR's commitment to the content determinacy constraint ought to be understood. The consideration in 4.2. might indeed provide a basis for *preferring* the latter understanding of the constraint in question. I contend that this response provides a fruitful question for further investigation.

5. Conclusions and Future Research Directions

The layout of this thesis was as follows. Chapter 1 consisted in an introduction to the research aims of the thesis, its methods, and the central research problem addressed (section 1.1), followed by an explanation of the theoretical significance of the topic and the justification for undertaking this project (1.2), and finally, a statement on the ethics involved (1.3). Chapter 2 consisted in explicating the broader theoretical backdrop of the notion of representation and its place and use within cognitive science (2.1. - 2.2.), and an explication of the *content determinacy constraint* for theories of representational content. Chapter 3 consisted in explicating DAMR in full. 3.1 consisted in an explication of the "computational theory proper" in articulation with DAMR 3.2. consisted in an explication of the two kinds of content — mathematical and cognitive — which DAMR maintains. 3.3. consisted in an explication of DAMR's realist construal of representational vehicles. Finally, 3.4. consisted in an explication of how DAMR purports to answer to the content determinant constraint outlined in chapter 2, section 2 (3.4.). Chapter 4 consisted in evaluating whether or not DAMR provides the basis to answer the content determinacy constraint. 4.1. consisted in explicating previous criticisms to DAMR. In section 4.2., I raised a theoretical concern which suggests that at least one of the pragmatic functions DAMR maintains might *not* provide the basis for the attribution of determinate content in some cases (and hence, might not meet the content determinacy constraint). 4.3. consisted in raising five potential responses to the concern raised in 4.2. and addressing those responses in turn. Chapter 5 consisted in presenting my conclusions — including answering the central research problem raised in this

thesis and, additionally, suggesting four considerations for potential future research. Chapter 6 comprises the references cited throughout the thesis.

In this thesis, I examined the deflationary account of mental representation (DAMR) proffered by philosopher Frances Egan. DAMR aims to provide an account of the nature and function of representation in cognitive neuroscience. To this end, I showed how DAMR maintains a realist construal of mental representational vehicles and a pragmatic account of mental representational content (Egan, 2010, 2014, 2017, 2018, 2020). I explicated how a realist construal of representational vehicles takes such vehicles to be actually-existing states or structures in a cognitive system, and how a pragmatic account of representational content takes such content to be ascribed by theorists for broadly *pragmatic* (or heuristic) reasons. I emphasized how, in particular, DAMR is motivated by the search to offer an explanation of representational content. Evaluating DAMR's explanation thereof thus formed the central investigation of my thesis.

The primary research problem that I aimed to address in my thesis was: does DAMR meet the *content determinacy constraint* on an account of representational content? To this end, I outlined the content determinacy constraint as stating that an account of representational content must provide the basis for the attribution of *determinate content* to a posited state or structure (Egan, 2020). I focused on one of two potential understandings of this constraint (outlined in chapter 2.2.). I narrowed my focus to one of the pragmatic functions which DAMR maintains, namely, the "placeholder" function (PF). I raised a theoretical concern which stated that PF might be vulnerable to content *indeterminacy*, and that such vulnerability might threaten to be generalizable or extended to other pragmatic functions DAMR maintains. This served to address the primary research problem stated above.

I thus conclude by answering the central problem posed in this thesis in the following way: there is reason to think that at least one of DAMR's pragmatic functions (namely, PF) might be vulnerable to content indeterminacy (as outlined in chapter 4.2.). There is, hence, reason to think DAMR does *not* meet the content determinacy constraint for every pragmatic function, in every case. However, it is this is a *limited* conclusion which requires qualification. It is limited in the following two ways. The first way in which it is limited is that the content indeterminacy concern raised in chapter 4 (4.2.) only applies to *PF* — not to all the pragmatic

functions DAMR maintains. Thus, the conclusion should *not* be understood as a suggestion that *DAMR* fails to meet the content determinacy constraint. I will suggest below that a fruitful research direction would be to investigate how PF *does*, in fact, interrelate with the other pragmatic functions DAMR maintains. A clear elucidation of this point could resolve the content indeterminacy concern.[47]

The second way in which it is limited is that there is more than one way of understanding the content determinacy constraint. I examined the concern I raised with regard to one such understanding of the constraint, namely, one which takes the constraint to be a constraint about the attribution of a certain kind of content (namely, determinate content). However, I did not examine the concern raised with regard to the other way of understanding the constraint on which the constraint is about *mere attribution* (see 2.2.). Thus, my conclusion might not apply to the content determinacy constraint which construes it in this latter sense. A clear elucidation of this point, too, could resole the content indeterminacy concern (as outlined in the fifth potential response raised in chapter 4.3.).

Finally, there are outstanding questions which remain. In particular, I want to highlight that there are multiple aspects of DAMR that could still use further investigation. There are four areas, in particular, which I want to highlight for future research investigation. The first of these is to clarify how we ought to understand the content determinacy constraint. In this thesis, I followed the assumption that the content determinacy constraint is a constraint about providing the basis for the attribution of a certain kind of content, namely, *determinate* content. However, as I outlined, Egan's own formulation of the constraint — and her subsequent discussion thereof — leave room for interpreting the constraint differently. Specifically, it leaves room for interpreting the constraint as being, centrally, about a theorists *attributive* practices. Clarifying how we ought to understand the way in which DAMR is committed to such a constraint is a fruitful area for further research.

The second future research area I want to highlight is this: it appears to be highly important to clarify how the various pragmatic functions DAMR maintains interrelate. This will serve to elucidate how exactly we are to understand how the pragmatic considerations specified. It will have the further effect of helping to clarify the additional question of how exactly such pragmatic functions ultimately provide the basis for the attribution of

[47] It could, for example, show that some other pragmatic factors limit PF in such a way that content indeterminacy of the kind raised is avoided.

determinate content in specific cases. Additionally, a clearer assessment of DAMR's various pragmatic functions might also serve to help further clarify the conclusion of this thesis: that it is unclear that DAMR meets the content determinacy constraint with regard to every pragmatic function it maintains

The third future research area I want to highlight is this: with the previous research areas in mind, another potential area of research would involve examining more closely the extent to which the content determinacy concern presented in 4.2. could be *generalized* or *extended to* other pragmatic functions DAMR maintains. Carrying out further investigation could serve to clarify if the worry presented affects other pragmatic functions. However, this is secondary to the previous two research areas identified. This is because either one of those research areas, if resolved, could thereby resolve the concern raised in 4.2. (as I noted above). This would have the effect of preventing a potential generalization or extension of the concern raised in 4.2.

Finally, the fourth future research area I want to highlight is this: the way in which — and to what extent — DAMR meets the variety of additional constraints on representational content and the additional, broader conditions on an account of mental representation is left open. I did not highlight these additional constraints in this thesis. This is because doing so was beyond its scope. However, the considerations in chapter 4 (4.2.) potentially bear on some of the constraints directly. For example, one of the additional constraints is the constraint which says that a representation must be able to *mis*represent.[48] Egan (2020) sees this constraint as *constitutively connected* to the content determinacy constraint. Thus, *if* it were to turn out that DAMR *is* vulnerable to content indeterminacy (for example, in PF), *then,* as a consequence, DAMR would so too be vulnerable to failing to underwrite *misrepresentation*, too. Further investigation into this — and the other content constraints and account conditions — hence provides another fruitful avenue for future research.

[48] I outline this notion briefly in chapter 2 (2.2.), footnote 12 and 21

6. References

Bechtel, William; Mitchell Herschbach (2010). “Philosophy of the Cognitives Sciences." In Fritz Allhoff (Ed.), *Philosophies of the Sciences*. Wiley-Blackwell. pp. 239-261

Betchel, William (2016). “Investigating Neural Representations: The Tale of Place Cells.” *Synthese* 193(5), Special Issue on Neuroscience and its Philosophy (2014-2015):1287-1321.

Chomsky, Noam (1995). “Language and Nature.” *Mind* 104(413):1-61.

Coelho Mollo, Dimitri (2017). “Content Pragmatism Defended.” *Topoi*, 39:103–113 (published 2020).

Cummins, R. (1989). *Meaning and Mental Representation*. Cambridge, MA: The MIT Press.

Dennett, Daniel (1978). “Toward a Cognitive Theory of Consciousness.” *Minnesota Studies in the Philosophy of Science* 9:201-228

Dennett, Daniel (1987). *The Intentional Stance*. Cambridge, MA: The MIT Press.

Dennett, Daniel (1991). *Consciousness Explained.* Boston: Little, Brown and Co.

Drayson, Zoe (2014). “The Personal/Subpersonal Distinction.” *Philosophy Compass*, 9(5):338-346.

Dretske, Fred (1981). *Knowledge and the Flow of Information*. Cambridge, MA.: The MIT Press.

Dretske, Fred (1986). “Misrepresentation.” In Bodgan R. (Ed.). *Belief: Form, Content and Function* (pp. 17-36). New York: Oxford University Press.

Egan, Frances (1995). “Computation and Content.” *The Philosophical Review*, 104(2):181-203.

Egan, Frances (1999). “In defense of narrow mindedness.” *Mind&Language*, 14(2), 177-194.

Egan, Frances (2009). “Is there a role for representational content in scientific psychology?” In Dominic Murphy; Michael A. Bishop (Eds.), *Stich and His Critics*, pp. 14-29 Wiley-Blackwell.

Egan, Frances (2010). “Computational models: a modest role for content.” *Studies in History and Philosophy of Science Part A,* 41(3):253-259.

Egan, Frances (2014). “How to think about mental content.” *Philosophical Studies*, 170 (1):115-135.

Egan, Frances (2017). “Function-Theoretic Explanation and the Search for Neural Mechanisms.” In *Explanation and Integration in Mind and Brain Science*, pp. 145-163. Oxford, UK.

Egan, Frances (2018). “The nature and function of content in computational models.” In Mark Sprevak & Matteo Colombo (Eds.), *The Routledge Handbook of the Computational Mind*, pp. 247-258. London, UK and New York, USA: Routledge.

Egan, Frances (2020). “A Deflationary Account of Mental Representation.” In J. Smortchkova, K. Dolega, T. Schlicht (Eds.), *What are Mental Representations?* (pp. 26-54), New York: Oxford University Press.

Fodor, Jerry (1975). *The Language of Thought*, Cambridge, MA.: Harvard University Press.

Fodor, Jerry (1981). *Representations: Philosophical Essays on the Foundations of Cognitive Science*. Cambridge, MA.: The MIT Press.

Fodor, Jerry (1990). *A Theory of Content and Other Essays*. Cambridge, MA.: The MIT Press.

Grush, Rick (1997). “The architecture of representation.” *Philosophical Psychology*, 10(1):5-23.

Gładziejewski, Paweł; Marcin Miłkowski (2017). “Structural representations: causally relevant and different from detectors.” *Biology & Philosophy*, 32:337-355.

Horgan, T.; Graham, G. (2012). “Phenomenal intentionality and content determinacy." In R. Schantz (Ed.), *Prospects for meaning* (pp. 321–344). Boston: De Gruyter.

Hurley, S. L. (1998). “Vehicles, Contents, Conceptual Structure, and Externalism.” *Analysis*, 58(1):1-6.

Hutto, Daniel D.; Eric Myin (2020). “Deflating Deflationism about Mental Representation.” In J. Smortchkova, K. Dolega, T. Schlicht (Eds.), *What are Mental Representations?* (pp. 79-100), New York: Oxford University Press.

Isaac, Alistair M. C. (2018). “Computational thought from Descartes to Lovelace.” In Mark Sprevak & Matteo Colombo (Eds.), *The Routledge Handbook of the Computational Mind*, pp. 9-22. London, UK and New York, USA: Routledge.

Lettvin, J. Y.; H. R. Maturana; W. S. McCulloch; W. H. Pitts (1959). “What the Frog’s Eye Tells the Frog’s Brain.” *Proceedings of the IRE*, 47(11):1940-1951.

Marr, David (1982). *Vision*. Henry Holt: New York.

Mendelovici, Angela. (2013). Reliable misrepresentation and tracking theories of mental representation. *Philosophical Studies*, 165(2):421-443.

Miller, George A. (2003). The cognitive revolution: a historical perspective. *TRENDS in Cognitive Science*, 7(3):141-144.

Millikan, Ruth G. (1984). *Language, Thought, and Other Biological Categories.* Cambridge, MA: The MIT Press.

Millikan, Ruth G. (1991). "Perceptual Content and Fregean Myth." *Mind*, C(400):439-459.

Millikan, Ruth G. (1993). "Content and vehicle." In N. Eilan, R. McCarthy, and B. Brewer (Eds.), *Spatial Representation*, pp. 256-68. Oxford: Blackwell.

Neander, Karen (2006). "Content for Cognitive Science." In Graham MacDonald and David Papineau (Eds.), *Teleosemantics: New Philosophical Essays*, pp. 167-194. New York: Oxford University Press.

Neander, Karen (2017). *A Mark of the Mental: A Defense of Informational Teleosemantics.* Cambridge, MA.: The MIT Press.

Newell, Allen; Herbert A. Simon (1976). "Computer science as empirical inquiry: symbols and search." *Communications of the ACM*, 19(3):113-126.

Piccinini, Gualtiero (2008). "Computation without representation." *Philosophical Studies*, 137(2):205-241.

Piccinini, Gualtiero (2020). *Neurocognitive Mechanisms: Explaining Biological Cognition*. Oxford University Press.

Pitt, D. (2018). "Mental Representation." In E. N. Zalta (Ed.), *The Stanford Encyclopedia of*

Philosophy. Winter 2018 ed. https://plato.stanford.edu/archives/win2018/entries/mental-representation/.

Port, Robert F.; Timothy van Gelder (Eds.) (1995). *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, USA, and London, UK: The MIT Press.

Putnam, Hillary (1967). "Psychophysical Predicates." In *Art, Mind, and Religion*, W. Capitan and D. Merrill (Eds.), Pittsburgh: University of Pittsburgh Press. (Reprinted as "The nature of mental states," in *Mind, Language and Reality*, *Philosophical Papers 2*, Cambridge: Cambridge University Press, 1975, pp. 429–440.)

Quine, W. V. O. (1960). *Word and Object.* Cambridge, MA: The MIT Press.

Ramsey, William (2007). *Representation Reconsidered*. Cambridge: Cambridge University Press.

Ramsey, William (2020). Defending Representation Realism. In J. Smortchkova, K. Dolega, T. Schlicht (Eds.), *What Are Mental Representations?* (pp. 55-78), New York: Oxford University Press.

Rescorla, Michael (2020). "Reifying Representations." In J. Smortchkova, K. Dolega, T. Schlicht (Eds.), *What Are Mental Representations?* (pp. 55-78), New York: Oxford University Press.

Rupert, Robert D. (2009). *Cognitive Systems and the Extended Mind*. New York: Oxford University Press.

Ryder, Dan (2004). "SINBAD Neurosemantics: a Theory of Mental Representation." *Mind & Language*, 19(2):211-240.

Searle, John R. (1980). “Minds, brains, and programs.” *The Behavioral and Brain Sciences*, 3:417-457.

Searle, John R. (1993). *The rediscovery of the mind*. Cambridge, MA: MIT Press.

Sellars, Wilfred (1962). “Philosophy and the Scientific Image of Man.” *Frontiers of Science and Philosophy*, Robert Colodny (Ed.), Pittsburgh: University of Pittsburgh Press, pp. 35-78. Reprinted in Science, Perception and Reality (1963).

Shagrir, Oron (2012), “Structural Representations and the Brain.” *British Journal for the Philosophy of Science*, 63:519-545.

Shea, Nicholas. (2018). *Representation in Cognitive Science*. New York: Oxford University Press.

Shea, Nicholas (2007). “Content and Its Vehicles in Connectionist Systems.” *Mind and Language*, 22(3):246-269.

Smolensky, Paul (1990). “Representation in Connectionist Networks.” *Intellectica*, 9-10:127-165.

Smortchkova, Joulia; Kryzysztof Dołęga; Tobias Schlicht (2018). “Introduction.” In J. Smortchkova, Kryzysztof Dołęga, T. Schlicht (Eds.), *What Are Mental Representations?* (pp. 1-25), New York: Oxford University Press.

Sprevak, M. (2013). “Fictionalism about neural representations.” *The Monist*, *96*(4):539-560.

Sprevak, Mark; Matteo Colombo (2018). “Introduction.” In Mark Sprevak & Matteo Colombo (Eds.), *The Routledge Handbook of the Computational Mind*, pp. 9-22. London, UK and New York, USA: Routledge.

Turing, Alan M. (1936). “On Computable Numbers, with an Application to the Entscheidungsproblem.” *Proceedings of the London Mathematical Society*, 42(1):230-265.

Turing, Alan M. (1950). “Computing Machinery and Intelligence.” *Mind*, 59:433-460.

Varela F. J.; E. Thompson; E. Rosch (1991). *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: The MIT Press.

Von Eckardt, Barbara (2012). “The representational theory of mind.” In Keith Frankish & William Ramsey (Eds.), *The Cambridge Handbook of Cognitive Science*. (pp. 29-49). Cambridge University Press.