

Reconsidering RepStat Rules in Dialectical Games

Simon Wells¹, Mark Snaith²

¹Edinburgh Napier University, 10 Colinton Road, Edinburgh, EH10 5DT, Scotland, UK

²Robert Gordon University, Garthdee House, Aberdeen, AB10 7QB, Scotland, UK

Abstract

Prohibition of repeated statements has benefits for the tractability and predictability of dialogues carried out by machines, but doesn't match the real world behaviour of people. This gap between human and machine behaviour leads to problems when formal dialectical systems are applied in conversational AI contexts. However, the problem of handling statement repetition gives insight into wider issues that stem partly from the historical focus on formal dialectics to the near exclusion of descriptive dialectics.

In this paper we consider the problem of balancing the needs of machines versus those of human participants through the consideration of both descriptive and formal dialectics integrated within a single overarching dialectical system. We describe how this approach can be supported through minimal extension of the Dialogue Game Description Language.

1. Introduction

In their recent chapter on Argumentation-based Dialogue [1], Black et al provide a nice overview of how argumentation can form a basis for dialogical interaction. The section on repetitive statements reproduces what could be described as the “standard treatment” of repetition in dialogue games. The simple solution, harking back to Hamblin, to handle an agent repeating the same assertion, and thus adversely affecting the termination guarantee, is to prohibit them from doing so. The basic idea is to prohibit utterance of any statement that is already in the speaker's commitment store or to disallow repetition of any move [2]. Whilst this is necessary in order to provide guarantees about how such agent dialogues proceed, and to make the resulting dialogues tractable, it does reduce the natural alignment between agent-based dialogues and human-based dialogues, one of which benefits from highly constrained rules, and the other of which is characterised by openness and flexibility.

People repeat statements during dialogue for many reasons, but there is good motivation to restrict such behaviour by software agents. So the question arises of how to balance the two, providing tractable protocols for software agents whilst simultaneously maintaining sufficient flexibility and realism for human participants in the context of conversational AI [3]? In the current work we consider conversational AI to be the application of AI techniques to dialogical interaction between people and machines. Whilst much recent work focuses upon the application of ML and other sub-symbolic techniques to this domain, there is still an important

CMNA'22: *International Workshop on Computational Models of Natural Argument*, September 12, 2022, Cardiff, UK

✉ s.wells@napier.ac.uk (S. Wells); m.snaith@rgu.ac.uk (M. Snaith)

🌐 <https://www.simonwells.org/> (S. Wells); <https://www3.rgu.ac.uk/dmstaff/snaith-mark/> (M. Snaith)

🆔 0000-0003-4512-7868 (S. Wells); 0000-0001-9979-9374 (M. Snaith)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

role to be played by symbolic approaches, dialectical games being a promising starting point from within the argumentation domain.

ML-based conversational AI has difficulty producing realistic extended dialogue, in particular, goal oriented, strategic dialogue in which the participants both set out their own positions, choosing what to say and what not to say at any given time, and also respond directly and opportunely to the utterances of each other, in order to satisfy both individual goals and joint aims. These are the sorts of dialogical interaction that are prevalent in many human to human dialogues, providing facets that underpin both the perception of naturalness and flexibility. However ML-based Conversational AI is capable of increasingly natural seeming, single or limited step interactions, for example those found in question-answer dialogues, but rapidly lose coherence and focus across longer exchanges. Conversely, structured dialogue approaches, such as those based on dialectical systems, can provide structure that enables goal-oriented, strategic interaction over extended sequences of utterances to occur. Consequently there is clearly still a role for dialectical systems within conversational AI.

In essence, we propose bringing together the well studied domain of formal dialectic with the less well studied domain of descriptive dialectic. Whilst ostensibly these are two distinct areas of concern, the former concerned with applying simple rules to generate dialogue that hold desirable characteristics, the latter begins with actual dialogue and attempts to identify the rules that capture such behaviour. Whilst study of formal dialectic exclusively is ideal for capturing highly structured dialogical interaction between people, it is more appropriately applied to dialogue between software, rather than human, agents, a context in which rigorous and tightly specified rules is a virtue. By bringing the two together in a cohesive and coherent way we get to provide a foci for software agent behaviour (where repetition is discouraged), and another for human-agent behaviour (where repetition is permitted when necessary, if not actively encouraged), within a context that allows smooth transition between the two. The resulting system is one that not only makes some advances in terms of dialectical systems as a basis for conversational AI, but provides a mechanism for resolving problems like RepStat which involve a need both to constrain behaviours generally whilst permitting them at other times and enabling the participants to make the choice about which to apply.

In the rest of this paper we delineate some aspects of the relationship between descriptive dialectic and formal dialectic, identify some ways in which co-occurring formal and descriptive dialogues give rise to new ways to monitor ongoing dialogues and to evaluate them at completion, and identify how some features of the Dialogue Game Description Language [4] can be reused to model concurrent formal and descriptive dialogues.

2. Related Work

Repetition in formal dialectical games has been studied frequently over the years and Black et al. [1] summarise the problem nicely, as a tension between the need, particularly within software agent dialogues, to ensure dialogue termination through a combination of clear rules, i.e. if an agent can't make a legal move or can't utter anything that it hasn't already conceded then it must concede and the dialogue terminates. Tactical aspects of repetition within formal dialectics were also studied by Parsons et al. [2].

The field of conversational AI has seen rapid expansion. A good recent, general overview is due to McTear [3]. While its roots can be traced back to primitive systems such as *ELIZA* [5], recent developments in artificial intelligence have revealed the potential for more sophisticated conversational interfaces with complex software systems. One of the most ubiquitous applications of conversational AI is chatbots [6]. These find uses across a range of domains, such as customer service [7], health care [8] and education [9]. Many such chatbots are based on natural language understanding (NLU) and natural language processing (NLP) techniques, through the extraction of intents, entities and contexts from statements made by human users [6]. In general, however, chatbots do not have an awareness of conversational structure and simply respond on the basis of the understood intent from the user. Conversational Agents, by contrast, do incorporate a degree of dialogue tracking and management [10]. This in turn allows for mixed-initiative dialogue, where an agent can take the lead in initiating a conversation [11].

3. Limits of RepStat Rules

Formal dialectic games assume a form of unwinding, where the acceptance of certain propositions and arguments leads to the acceptance of previous propositions. For example, in *PPD* [12], there is a structural rule that requires a participant to be committed to an argument's conclusion (C) if they are also committed to the propositions and rules that lead to that conclusion. Such commitment arises through the dialectical machinery and could happen quite far from the proponent's original claim of C. This mismatches with expected real-world behaviour where the proponent of C might be expected to restate it, either to confirm their opponent's acceptance, or to reinforce the defeat.

Repetition in dialogue can also be used to enforce quality while also having a strong rhetorical effect. Consider the following real example: in 1997, BBC interviewer Jeremy Paxman asked then UK Cabinet Minister Michael Howard "*Did you threaten to overrule him?*" no fewer than twelve times in the face of ostensibly unacceptable answers¹. Repeating the same question so many times reinforced the perceived lack of answer in a way that an alternative move (e.g. different instantiations of "*that doesn't answer the question*") might not.

One solution to dealing with the limitations of existing RepStat rules might be to investigate all of the circumstances in which repetition might be legal, and to formulate permissive rules that are exceptions to the general prohibition of repetition. A complementary approach might be to blanket permit repetition, but to identify all of the circumstances when it shouldn't be allowed. Or perhaps some hybrid of these approaches. However, whilst both approaches are firmly within the arena of topics that should continue to be investigated as topics within formal dialectical research, they still retain some troubling aspects, firstly, that people and machines are treated as equal participants, and secondly, that listing circumstances by extension, is a perilously fragile approach that risks the same criticisms that have traditionally been applied to symbolic AI, i.e. that as soon as the scenario doesn't quite match the model, the model fails to be usefully applicable.

This issue isn't restricted solely to the RepStat problem, it is a general criticism of formal dialectics more generally when applied to managing real-world dialogue. Whilst formal dialectic

¹The exchange can be viewed here: <https://www.youtube.com/watch?v=Uwlsd8RAoqI>

tics are analytically excellent tools, and incredibly useful for inter-agent communication, their utility in human-human, human-agent, or mixed initiative contexts is reduced. The problem is that formal dialectics captures the details of specific human interactions but elides many of the generalities of those same interactions.

However, descriptive dialectics are the missing part of the problem. They are meant to capture, describe, and model exactly those circumstances in dialogue in which formal dialectics are less accomplished. Descriptive dialectics provide the missing piece. Whereas formal dialectics capture dialogical behaviours that are carefully circumscribed and focused, descriptive dialectics capture the larger, more general context in which people engage in dialogue.

That formal dialectics have been focused upon, to the near exclusion of descriptive dialectics, is completely understandable. From a computational viewpoint, the precision of formal dialectical rules gives both explicit and predictable structure, but importantly also sets boundaries upon what can be said which in turn can help to make inter-agent communication more predictable and thus more computationally tractable. One solution is to bring descriptive dialectics to the fore, to recognise that the two approaches, formal and descriptive, address complementary aspects of dialogue analysis and modelling.

To summarise, dealing with the RepStat mismatch is a general instance of a wider issue that comes into sharp focus as we begin to look toward dialectical systems as a tool to achieve conversational AI. This issue is the tension between the constraint inherent within formal systems, designed to model specific problems within dialogue and to enable the exploration of the dialogues reachable from simple but perhaps unrealistic rules, and the liberty inherent to descriptive systems, designed to describe real world phenomena.

4. A New Approach to Handling Repetition

Having previously established how the problem of repetition is related to an issue of scoping between formal and descriptive approaches, in this section, we now show how some existing features of the DGDL can be repurposed to support descriptive and formal dialectical games running concurrently.

The DGDL is a tool for describing dialectical games which, together with a DGDL run-time environment or *execution engine* such as the Dialogue Game Execution Platform (DGEP) [13] or ADAMANT [14], enable executable dialectical game to be specified and run. This approach enables flexibility in terms of the specific rules that any given game comprises. DGDL games are specified according to an Extended Backus-Naur Form grammar. Briefly, a DGDL system comprises one or more named games. An individual game has a composition which defines the participants, roles that can be assigned to the participants, a turn taking structure, and the artefact stores associated with the game. In addition a game has a set of rules, regulations that are applied at dialogue commencement, after each move, or after each turn, regardless of the move made by current player, and a set of interactions, which defines the specific moves that players are permitted, mandated, or prohibited to make.

A DGDL execution engine does three main jobs, firstly, maintaining dialogue state, secondly, determining the set of legal moves, and finally, verifying whether a given move conforms with or otherwise, the set of legal moves. Dialogue state corresponds directly to the constituents

		Initial Situation		
		Conflict	Open Problem	Unsatisfactory Spread of Information
Main Goal	Stable Agreement/Resolution	Persuasion	Inquiry	Information Seeking
	Practical Settlement/Decision (Not) to Act	Negotiation	Deliberation	
	Reaching a (Provisional) Accomodation	Eristic		

Table 1
Walton & Krabbe's systematic survey of Dialogue Types [12, pp. 80]

of the clause associated with the *Composition* keyword, essentially the game pieces and board. Determining the set of legal moves and whether a given played moves is legal depends upon the transcript, the set of previously played moves, the current state of the games pieces and board, and the set of moves available in the current game.

Much research has focused upon formal dialectical games where each game is based around a particular dialogue type. In this approach, dialogue types are generally based upon the influential typology due to Walton & Krabbe [12]. In their typology, Walton & Krabbe identify a number of stereotypical dialogue types, originally six distinct elemental dialogue types as illustrated in Table 1. These include persuasion, negotiation, inquiry, deliberation, information seeking, and eristics, as well as a seventh *mixed* or compound type that includes elements from the aforementioned elemental dialogues. In Walton & Krabbe's approach, dialogue types are distinguished based upon the following criteria: initial situation, individual goals, and main/joint goals. Whilst dialogue types have influenced the foci of much dialectical game research and thus the kinds of games that the DGDL can support, another innovation from Walton & Krabbe, the concept of shifts and embeddings has also influenced how the DGDL supports relationships between arbitrary games. A dialogue shift occurs when the participants in a dialogue move from an instance of one dialogue type to an instance of a new dialogue of the same or a different type. For example, shifting from persuasion to negotiation as occurs during the fallacy of bargaining [15]. An embedding occurs when the shift is symmetrical, moving to a child dialogue and then returning to the parent dialogue upon completion.

The DGDL *system* keyword was initially designed to support shifts and embeddings of the types initially identified by Walton & Krabbe. A DGDL system is a group of one or more dialectical systems that can operate collectively enabling players to shift from an instance of one dialogue game to an instance of another game. Note that the term embedding here is merely descriptive, to conjure up the situation where a sub-dialogue is entirely symmetrically encapsulated within a parent dialogue. There is also no limitation on the number of sub-dialogues that can be embedded within each other, just that the term embedding is a terminological shortcut to identify that the dialogue returns to the parent rather like a computational stack being pushed and popped. However dialogue shifts can also occur linearly, without embeddings, involving serial shifts from one dialogue to the next, and so on.

The mechanism for effecting a shift can be distinct, involving a move that causes a shift to

occur immediately, or gradual, where the current state of the dialogue licenses a shift. Licensing occurs when the set of legal moves for the current dialogue overlaps with the set of moves from one or more destination dialogues. The player then plays their next move which can be from one of three categories of legal response. If the responsive move is solely from the subset belonging to the destination dialogue then the shift occurs successfully but if the move is from the subset belonging solely to the source dialogue, then the shift fails to occur. In some cases the move can belong to both source and destination dialogues, resulting in an indeterminate state which will last until some subsequent move causes it wholly within the source or the destination dialogue. This indeterminate state may last as long as necessary to successfully complete the shift or to return to the source dialogue, and corresponds to what Walton & Krabbe refer to as a *glissando* style shift, in which the shift is not abrupt, but instead extends over as many moves as are necessary.

Whilst couched as a domain specific language (DSL) for describing dialectical games, the DGDL is more properly described as a DSL for describing dialectical systems. The distinction is important here because a dialectical system, in Hamblin's description [16, pp. 255], isn't restricted merely to formal games but can also incorporate descriptive games. The DGDL fully implements Hamblin's conception of "dialectical systems being pursued *descriptively*, or *formally*" where "neither approach is of any importance on its own" [16, pp. 256] to formally support both between formal and descriptive games. In one sense, this is quite straightforward to support using existing DGDL features, a unified dialectical system is a DGDL description that supports at least two games, where one game is a formal game, and the other game is a descriptive game. Furthermore, to be effective, the games should support shifts, back and forth as necessary, between them.

However this is not sufficient to effectively handle the RepStat problem which really needs both games to run concurrently. In this way the descriptive game describes an outer, more permissive game that supports more generalised and human-like interactions, in which repetition is permitted. The formal game meanwhile represents an inner, more restricted game that supports more constrained and computational tractable behaviour. To some degree the difference between formal and descriptive, or inner and outer, or permissive and restricted, games, is analogous to the difference between the Permissive Persuasion Dialogue (PPD) and Rigorous Persuasion Dialogue (RPD) games of Walton & Krabbe [12]. By running concurrent formal and descriptive games, there are resultant move sets available separately to the human and machine participants that play to the respective strengths of each.

There are three additional points to consider. Firstly, legal move determination for concurrent games works identically to that for shifts. Secondly, the DGDL doesn't specify what a player should do, only what they can legally do. It is therefore down to the individual participants to decide whether to select a move from the formal set or the descriptive set. It is however expected that machines should be designed to select primarily from the formal set, in order to maintain the computational advantages of formal games, whilst human participants are free to select any move from either the descriptive or the formal set. Note that a machine player might, under some circumstances elect to make a move from the descriptive set, but it is expected that the proportions of such moves is low. Note that the main goal of running concurrent formal and descriptive games is to provide the flexibility necessary to support more natural engagement for the human participants whilst minimising any resultant impact on the machine participants.

Thirdly, this has an impact on the process of designing dialectical games. To exploit the DGDL approach to dialogue shifts requires careful design of the constituent games that make up a DGDL system so that they work together. Similarly, to run concurrent formal and descriptive games requires that the elements of each game's composition significantly overlap so that the effects of a move within one game affect the game board of the other, and vice versa.

To summarise, by using existing DGDL features, with slight execution environment modifications, and by designing new games to take advantage of the unified, concurrent approach, new, more flexible, dialectical systems can be developed that account for the competing needs of computational tractability and human-friendly flexibility necessary for human-friendly conversational AI interfaces.

5. Motivating Examples

Here, we motivate our proposed approach with examples taken from a health care domain. Development of technological solutions to support health care has gained significant traction in recent years, especially in the area of behaviour change for chronic illness. The use of coaching-based Behaviour Change Support Systems (BCSS) allows medical practitioners to provide patient interventions towards changes in areas including diet, exercise, and social support [17]. The use of conversational AI allows patients to interact with such systems in an engaging way that simulates discussions they might have with a real health care professional [18, 19, 20]. Underpinning these conversational interactions with models of dialogue [21] can help address potential ethical issues around trust and explainability [22].

One area of health coaching that particularly benefits from structure is Motivational Interviewing (MI), a technique that helps patients explore their own motivations for change [23, 24]. This structure however needs to be flexible, with the ability to adapt and react to a patient's self-exploration. It is also essential that repetition be allowed as a way of summarising what has been said, and ensuring understanding of any agreed goals or strategies.

Our examples are taken from the *Patient Consultation Corpus* (PCC) [25], consisting of simulated discussions between real health care professionals and an actor portraying a patient with a defined persona. Three discussions in the PCC involved a Motivational Interviewer, and all contain examples of exchanges that would be difficult to fully model with formal dialectic. Our intention with these examples is not to provide a full dialectical analysis of the exchanges, but rather highlight some specific phenomena that illustrate the need for the approach described.

5.1. Example 1 - repetition as confirmation

This example is taken from [25, D2.C1]. In this session, a motivational interviewer and a dietitian attempt to coach the patient to eat more healthily. Towards the end of the session, the motivational interviewer makes the following statement:

So, just to pull together what we've talked about and the things that you've said you feel could do to try to maintain your weight...

He then goes on to restate what was agreed to. This restatement of the agreement is a vital step in MI to ensure there is mutual understanding of the patient's goals. From a formal

perspective, it could be assumed (based on the rules of a hypothetical game) that the original statement of the goals leads to mutual commitment and thus agreement. While such an approach might be suitable for communication exclusively between software agents, it does not reflect natural, inter-human dialogue. First, when coming to an agreement it is important that all parties involved know they are agreeing to the same thing. Second, revisiting elements of an agreement is common practice, so failing to incorporate this step would lose an aspect of natural communication.

5.2. Example 2 - repetition as clarification

This example is taken from [25, D2.C2]. In this session, the patient has received a lot of information and suggestions at once and is feeling overwhelmed. One of the health practitioners makes the following suggestion:

I think, picking up from what Colin was saying, it's easy to get overwhelmed with all of these different things, that someone's told you...And I think, what Colin was just talking about there, was really good, about breaking things down into small chunks.

This then triggers a continuation of the dialogue where some of the previously mentioned actions were revisited. This is again an example of restating being essential as a consequence of natural inter-human interaction. The previous utterances would correspond to commitment, at least on the part of the health practitioners, and thus not afford the opportunity to restate them in a more accessible way that might lead to agreement from the patient.

5.3. Example 3 - deviation to resolve impasse

This example is taken from [25, D2.C3]. In this session, a motivational interviewer and a dietitian are struggling to convince a patient of the benefits of a healthy diet. The patient has used *Google* to conduct their own research, and what they have found disagrees with what they are being told. In an attempt to resolve the impasse, the motivational interviewer breaks the flow of dialogue with the following question:

Okay. So, I get that you have researched. Can I just ask, in your work, if you research a piece for your work, what do you search through?

This leads to a discussion around the skillset required for the patient to do her job, and whether or not it could be done using *Google*. The aim of this exchange is to have the patient reflect on their over-reliance on what the internet tells them. However, this exchange is not a part of the core dialogue; rather, it is an interruption designed to adapt to the challenging circumstances the health practitioners are facing. In a formal dialogue game, this would not be permitted, whereas in a descriptive game - which allows for strategic interruptions - it would.

6. Discussion

In this paper we've described how a reuse of existing DGDL machinery, namely the *system* keyword and related infrastructure, can allow both descriptive and formal dialectical games

to be represented and to run concurrently. This has been applied not only to develop *unified* dialectical systems, integrating both dialectical and formal aspects, but to investigate how to solve problems characterised by the need to maintain tractability for software agent participants, whilst also maintaining flexibility and naturalness for human participants. This has been proposed in the context of the new frontiers in conversational AI, a timely endeavour, within which dialectical games should play an important role.

There remain many open questions however, for example, how do descriptive and formal games differ in terms of the kinds of rules that constitute them. We've assumed herein that both types of game operate in the same space and are constructed from the same basic kinds of rules, reusing those rules already extant in the DGDL. However we conjecture that descriptive games, in general, will be formulated more frequently from a combination of permissive and prohibitive rules, whereas formal games are generally prescriptive and prohibitive. This reflects the intuition that descriptive games capture *can (not) do* constraints whereas formal games capture *must (not) do* constraints.

The study of how formal and descriptive dialectical games differ, in a formal sense, will be addressed in future work as the programme of descriptive analysis of real world dialogue develops. This might lead, if necessary, to additional DGDL rules to better account for both real world, and ideal, dialogues, as captured respectively by descriptive and formal dialectical approaches.

The next open question relates to the range of modalities associated with move selection by participants during a dialogue. For example, if two, or more, games are simultaneously active and represent poles of a descriptive-formal continuum, then this suggests new ways of instrumenting, tracking, and evaluating what is happening during an active dialogue. For example, a dialogue that proceeds mostly or entirely according to formal rules might be considered to be tighter or more focused, whereas one that proceeds more according to descriptive rules might be more natural and realistic. The determination of which is better, and when, according to different dialogical contexts, is yet to be determined.

Future work will concentrate upon the challenge associated with interleaving the rules from different, concurrently running games, investigation of the effect that this approach has upon determining the set of legal moves, and any resultant effect upon the tactical and strategic positions of the participants as a consequence.

In summary, this paper has, within the wider context of seeking to solve the *RepStat* problem, introduced a new approach to considering how to build better conversational AI systems.

References

- [1] E. Black, N. Maudet, S. Parsons, *Argumentation-based Dialogue*, volume 2, College Publications, 2021.
- [2] S. Parsons, M. Wooldridge, L. Amgoud, Properties and complexity of some formal inter-agent dialogues, *Journal of Logic and Computation* 13 (2003) 347–376. URL: <http://www.sci.brooklyn.cuny.edu/~parsons/publications/journals/jlc3.html>.
- [3] M. McTear, *Conversational AI: Dialogue systems, conversational agents, and chatbots*,

Synthesis Lectures on Human Language Technologies 13 (2020) 1–251. doi:10.2200/S01060ED1V01Y202010HHLT048.

- [4] S. Wells, C. Reed, A domain specific language for describing diverse systems of dialogue, *Journal of Applied Logic* 10 (2012) 309–329.
- [5] J. Weizenbaum, Eliza—a computer program for the study of natural language communication between man and machine, *Communications of the ACM* 9 (1966) 36–45.
- [6] E. Adamopoulou, L. Moussiades, Chatbots: History, technology, and applications, *Machine Learning with Applications* 2 (2020) 100006.
- [7] M. Adam, M. Wessel, A. Benlian, AI-based chatbots in customer service and their effects on user compliance, *Electronic Markets* 31 (2021) 427–445.
- [8] F. Amato, S. Marrone, V. Moscato, G. Piantadosi, A. Picariello, C. Sansone, Chatbots meet eHealth: Automating healthcare., in: *WIAIAH@ AI* IA*, 2017, pp. 40–49.
- [9] B. Heller, M. Proctor, D. Mah, L. Jewell, B. Cheung, Freudbot: An investigation of chatbot technology in distance education, in: *EdMedia+ innovate learning*, Association for the Advancement of Computing in Education (AACE), 2005, pp. 3913–3918.
- [10] A. Rastogi, X. Zang, S. Sunkara, R. Gupta, P. Khaitan, Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2020, pp. 8689–8696.
- [11] M. Snaith, J. Lawrence, C. Reed, Mixed initiative argument in public deliberation, *Online Deliberation* 2 (2010).
- [12] D. Walton, E. C. Krabbe, *Commitment in dialogue: Basic concepts of interpersonal reasoning*, SUNY press, 1995.
- [13] J. Lawrence, M. Snaith, B. Konat, K. Budzynska, C. Reed, Debating technology for dialogical argument: Sensemaking, engagement, and analytics, *ACM Transactions on Internet Technology (TOIT)* 17 (2017) 1–23.
- [14] S. Wells, The open argumentation platform (OAPL), in: *Proceedings of Computational Models of Argument. (COMMA 2020)*, *Frontiers in Artificial Intelligence*, IOS Press, 2020, pp. 465–476.
- [15] S. Wells, C. Reed, Knowing when to bargain, in: *Proceedings of the 1st Conference on Computational Models of Argument. (COMMA 2006)*, 2006.
- [16] C. L. Hamblin, *Fallacies*, Methuen and Co. Ltd, 1970.
- [17] B. A. Rosser, K. E. Vowles, E. Keogh, C. Eccleston, G. A. Mountain, Technologically-assisted behaviour change: A systematic review of studies of novel technologies for the management of chronic illness, *Journal of telemedicine and telecare* 15 (2009) 327–338.
- [18] N. Stein, K. Brooks, et al., A fully automated conversational artificial intelligence for weight loss: Longitudinal observational study among overweight and obese adults, *JMIR diabetes* 2 (2017) e8590.
- [19] R. B. Kantharaju, A. Pease, D. Reidsma, C. Pelachaud, M. Snaith, M. Bruijnes, R. Klaassen, T. Beinema, G. Huizing, D. Simonetti, et al., Integrating argumentation with social conversation between multiple virtual coaches, in: *Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents*, 2019, pp. 203–205.
- [20] A. Fadhil, Y. Wang, H. Reiterer, Assistive conversational agent for health coaching: A validation study, *Methods of information in medicine* 58 (2019) 009–023.
- [21] M. Snaith, D. De Franco, T. Beinema, H. O. Den Akker, A. Pease, A dialogue game for multi-

- party goal-setting in health coaching, in: 7th International Conference on Computational Models of Argument, COMMA 2018, IOS Press, 2018, pp. 337–344.
- [22] M. Snaith, R. Ø. Nielsen, S. R. Kotnis, A. Pease, Ethical challenges in argumentation and dialogue in a healthcare context, *Argument & Computation* 12 (2021) 249–264.
- [23] J. Hettema, J. Steele, W. R. Miller, Motivational interviewing, *Annu. Rev. Clin. Psychol.* 1 (2005) 91–111.
- [24] W. R. Miller, G. S. Rose, Toward a theory of motivational interviewing., *American psychologist* 64 (2009) 527.
- [25] M. Snaith, N. Conway, T. Beinema, D. De Franco, A. Pease, R. Kantharaju, M. Janier, G. Huizing, C. Pelachaud, et al., A multimodal corpus of simulated consultations between a patient and multiple healthcare professionals, *Language resources and evaluation* 55 (2021) 1077–1092.