Towards a Logic-Based Approach for Multi-Modal Fusion and Decision Making during Motor Rehabilitation Sessions

Fabio Aurelio D'Asaro
CRDC Tecnologie
University of Naples Federico II
Napoli, Italy

Antonio Origlia
URBANECO
University of Naples Federico II
Napoli, Italy
antonio.origlia@unina.it

Silvia Rossi
DIETI
University of Naples Federico II
Napoli, Italy
silvia.rossi@unina.it

***Abstract*—We introduce a general approach which aims at combining machine learning and logic-based techniques in order to model its user's cognitive and motor abilities. In the context of motor rehabilitation, hybrid systems are a convenient option as they allow both for the representation of formal constraints needed to implement a clinically valid exercise, and for the statistical modelling of intrinsically noisy data sources. Moreover, logic-based systems offer a transparent way to look at the decisions taken by an automated system. This is particularly useful when an AI system needs to interact with a therapist in order to assist therapeutic intervention, e.g. by explaining why a given decision is sound. This methodology is currently being developed within the context of the AVATEA project.**

***Index Terms*—Multimodal Fusion, Epistemic Probabilistic Event Calculus, Motor Rehabilitation**

I. INTRODUCTION

In this work, we introduce some of the ongoing activities in the context of the AVATEA project (Advanced Virtual Adaptive Technologies e-hEAlth). The project aims at developing an intelligent system to support the rehabilitation process of children with neuro-motor disorders. More specifically, AVATEA aims at creating an integrated system consisting of: (i) an adjustable seat, (ii) different types of sensors, and (iii) an interactive visual interface to perform rehabilitation exercises in the form of games. Such games/exercises are going to be specifically targeted at supporting therapeutic sessions for Development Coordination Disorders (DCD).

Although a significant amount of work has been done in the general area of motor rehabilitation with promising results [1], there is still a need for developing personalised therapeutic scenarios. Adaptation techniques typically only focus on maximising effort during the rehabilitation session [2]. However, it is also necessary to take into account parameters such as the individual subject's capabilities [3] and the child's emotional response, e.g. in terms of motivation and engagement [2].

In this direction, AVATEA aims to assist the activity of a therapist through the use of data acquired from its sensors. Machine learning techniques can process this data to profile the user's motor abilities, his/her psychophysiological state, and to monitor the child's response to the exercise being performed. The resulting user model can then be used e.g. to monitor the user's performance and decide what is the most appropriate rehabilitation strategy.

However, handling data coming from different sources requires a complex system able to integrate them and take decisions accordingly, i.e. a *multimodal system* [4]. Moreover, in existing rehabilitation games, the patient motivational state has been considered to evaluate the game effectiveness [5], without providing the possibility of taking decisions with respect to the online adaptation of the rehabilitation exercise.

Alongside machine learning techniques, we also intend to employ a logic-based system. Indeed, in a complex domain such as that of AVATEA, much of the experts' knowledge would have to be re-learnt from scratch (thus requiring a considerable amount of data) if we were to use machine learning techniques exclusively. Logic-based systems, on the other hand, are progressively becoming able to handle uncertain knowledge (e.g., using Probability Theory or Fuzzy Logic). This provides the opportunity to retain important parts of this knowledge even when it comes with a degree of uncertainty. Moreover, logic-based systems offer a transparent way to look at the information in AI systems. For example, a therapist might want not only to see *what* decisions were taken by the system but also *why* they were taken. A logic-based system is able to reconstruct the rationale behind the decisions taken by considering the chain of rules that were applied starting from the facts in the knowledge base. Such an advantage does not also apply to machine learning algorithms, that generally cannot provide explanations in human-readable terms. It is also worth noting that these systems can be used by an expert to sketch the causal relationships of a domain, and then use other techniques to learn the appropriate parameters when they are not available to the expert.

II. BACKGROUND AND RELATED WORK

Gamification strategies have proven to be extremely successful to engage young children in diagnostic or therapeutic exercises, even before the advent of digital gaming. By leveraging on Self-Determination Theory [6], the concept of *intrinsic motivation* has been applied to engage children in activities designed to provide therapists with reports about their competence levels in either cognitive or physical tasks. While games to test cognitive capabilities (see e.g. [7]) do not form a sharply defined class, games designed to test and improve motor skills are usually referred to as *exergames*. The effects of exergames have been found to be generally positive [8]. Therefore, combining Computerised Adaptive Tests (CATs) with gamification strategies results in systems that are able to engage young users in playful activities, while adapting the current challenge according to level of user competence. Furthermore, sessions are typically logged in order to provide detailed feedback to therapists. On the cognitive side, these adaptive systems have been designed to evaluate subjective well-being [9] and phonological acquisition [10] among others. Adaptive exergames have been used e.g. in the context of children with spinal impairments [11], and to test gross motor skills [12]. These works, however, typically use a single modality to implement CATs and do not consider social feedback as a part of the adaptation process to recover engagement. Considering the challenge posed both by multimodal fusion and by adaptation strategies, an integrated system for both fusion and decision making represents an ambitious goal, with potentially broad impact on the field of adaptive diagnosis and treatment of both cognitive and physical impairments.

Moreover, in the specific case of CATs, the domain knowledge is known to the developers, as the experimental procedure must remain safe and informative. It is necessary to keep records of the decisions taken by the system, in order to reconstruct and explain how the session was managed by the system. In this specific situation, statistical modelling alone is not advantageous because: (i) it would require a lot of data to *discover* elements that are already known, and (ii) it would make it difficult to provide human-readable feedback to the therapists.

In this context, hybrid systems are a convenient option as they allow both for the representation of formal constraints needed to implement a clinically valid exercise, and for the statistical modelling of intrinsically noisy data sources. Recently, logic-based approaches have been successfully applied to several fields of Artificial Intelligence, including (but not limited to): event recognition from security cameras [13], [14], robot location estimation [15], understanding of tenses [16] and natural language processing [17]. Due to the increasing relevance of Machine Learning and Probability Theory in AI, these frameworks and languages have gradually started employing probabilistic semantics (see e.g. [18]) to incorporate and deal with uncertainty. This has given birth to the field of Probabilistic Logic Programming (see e.g. [19]). For example [15], which is based on the *Situation Calculus* ontology, can model imperfect sensors and effectors. The Situation Calculus' branching structure makes these frameworks mostly suitable for planning under partial states of information. On the other hand, MLN-EC and ProbEC [13], [14] extend the semantics of the *Event Calculus* using Markov Logic Networks [20] and ProbLog [21] to perform event recognition from security cameras. In the proposed case study, the logical part of the architecture receives time-stamped events as inputs and processes them in order to detect complex long-term activities (e.g., detect that two people are fighting from the fact that they have been close to each other and moving abruptly during the last few seconds). Given their semi-probabilistic nature, these frameworks are able to handle uncertainty in the input events (ProbEC) or in the causal rules linking events and fluents (MLN-EC). We envisage that similar systems, especially the Event Calculus-based ones, could be employed as a way to perform fusion between different modalities.

III. THE AVATEA ARCHITECTURE

The proposed architecture is essentially a multimodal system. These were first formally defined in [4] as systems that "[. . .] process two or more combined user input modes such as speech, pen, touch, manual gestures, gaze, and head and body movements in a coordinated manner with multimedia system output". The possibility to handle multiple communication channels is expected to simplify interaction with the user and to result in a more natural way to control an automated system. Available modalities may be used in an *exclusive* or *concurrent* way, with no integration between them [22]. However, it is more often the case that a multimodal system processes multiple channels in a parallel and integrated way [23]. Moreover, whether it is more advantageous to adopt an *early* or *late* model for data fusion strongly depends on the amount of available knowledge about the domain. From a system design perspective, it is better to develop separate, more specialised, approaches to analyse each single data source and then fuse the results using a subsequent layer. However, when the domain knowledge is limited, key interactions among input modalities may be overlooked: in this case, early fusion is more adequate. Generally, the problem of deciding when to apply fusion is one of the main issues when designing multimodal systems (see e.g. [24]). In our domain, the amount of available knowledge about training exercises supports the adoption of a late fusion approach.

Figure 1 shows the envisaged architecture for the AVATEA project which main modules we will discuss in the following paragraphs.

A. Sensors

We are going to use different types of sensors, including: (i) pressure sensors, (ii) 2D and 3D cameras, (iii) motion detectors, and (iv) an EEG sensor. Pressure sensors, motion detectors and cameras are going to be used to make sure that exercises are being executed correctly by detecting front and back posture, head-pose, movement speed, balance and feet relative position. In some cases, the cameras may be used to let the user interact with the system, e.g. by pointing at the screen. Moreover, data from cameras and EEG data are going to help checking the user's current level of stress and engagement and e.g. decide whether the difficulty level of the exercise should be changed.

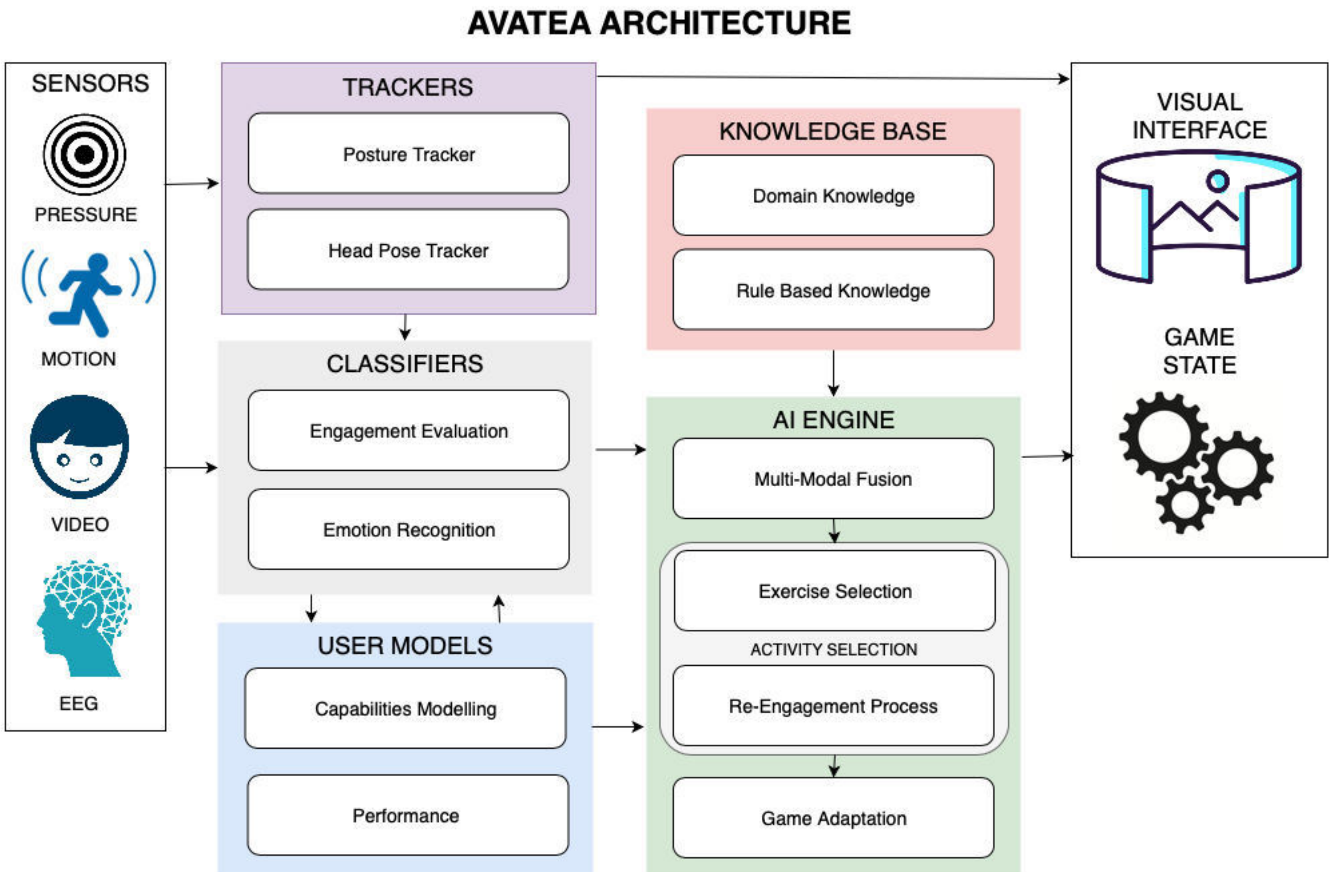

Fig. 1. The AVATEA architecture

B. Input Trackers

Adaptive games require the ability to dynamically track children movements. For example, movements must be taken into account if we want the system to automatically adapt the game speed to the children' physical capabilities. To this aim, pressure sensors, head pose and skeleton data from video images will be processed and then given on input to the game. Particular emphasis will be put on tracking children' posture: for instance, the head-pose will be a triggering event for some games. We are going to employ one of the several available skeleton detection algorithms. Existing methods include those using R-GBD or 2D cameras (e.g. the OpenPose library [25]) which can identify various positions, even those of ambiguous interpretation (sitting, three-quarter backward perspective, etc..).

C. Classifiers

In AVATEA, the focus is on human activities. These are very difficult to classify due to the diversity of individual conditions. Leveraging on the expressive power of deep networks as feature extractors, and by exploiting features modelling techniques of the human body, we will research and design novel algorithms for Social Signal Processing [26]. Video and EEG data will be used to monitor the attention of the children during the exercise, as well as his/her emotional state and engagement. The Modalities Recognisers classify the features extracted from the sensor data, and create a list of possible interpretations (N-Hypothesis) for the AI engine.

D. User Models

We will use personalised machine learning techniques to learn a model of the children abilities and interaction preferences [27]. In turn, these profiles will make the system able to recognise anomalies with respect to such model [28] so that it is possible to track improvements in the user's performance. Moreover, user's performance over time will be correlated with how they felt about the exercise (or similar exercises) in the past. This can be used to create a personalised exercise plan (e.g., exercise type, modality of execution of the exercise, etc).

E. AI Engine

The system we envisage implies the use of noisy data sources coming from multiple sensors and trained classifiers. These drive a decisions layer conducting adaptive rehabilitation exercises. If, on one hand, this makes it necessary to handle classification estimates with probabilistic reasoning, on the other hand, one needs to keep a rule-based structure to ensure clinical effectiveness and human-readable session summarisation. Hybrid systems, typically consisting of probabilistic rules, combine the best of the two approaches by allowing the definition of strict rules. These rules can be used to model the structure of the clinical procedure, and adapt to the information from the classifiers (i.e. confidence and probability distributions over classes). A user model based on probabilistic estimates can be used by a rule system to estimate the best course of action using expert knowledge encoded in a rule system. This user model can be then further processed in order to customise the therapeutic intervention, and therefore to raise the quality of the children experience, by also taking into account the behaviour of the child during the rehabilitation process. Indeed, the visual interface will be used to offer children the exercises as part of recreational activities that make use of detected social signals. This will offer a rehabilitation process based on games whose behaviour automatically adapts to the child.

1) EPEC: We propose the use of language *EPEC* (short for *Epistemic Probabilistic Event Calculus*) as a foundation for our methodology. Similarly to MLN-EC and ProbEC, EPEC is a language in the style of the Event Calculus for reasoning about actions, but goes beyond these languages in that it allows for the modelling of noisy sensors. Its foundations were laid in [29], and it has since then been extended in [30] to also include sensing actions and propositions conditioned on belief. We briefly introduce its syntax in the following.

In the tradition of reasoning about action languages, EPEC models a given domain using *fluents* (which represent properties of the world), *instants* (which represent time points at which events may occur) and *actions* (which represent actions under the control of the agent being modelled or the environment itself). The causal interactions between fluents and actions are captured by the specialised propositions below:

- the *v-proposition*

  $$F\ \textbf{takes-values}\ \langle V_1, \ldots, V_n \rangle$$

  states that fluent F can take values $V_1, \ldots, V_n$.
- the *i-proposition*

  $$\textbf{initially-one-of}\ \{(\psi_1, P_1), \ldots, (\psi_n, P_n)\}$$

  states that the environment is initially in one of the states $\psi_1, \ldots, \psi_n$ with probabilities $P_1, \ldots, P_n$.
- the *c-proposition*

  $$\theta\ \textbf{causes-one-of}\ \{(\psi_1, P_1), \ldots, (\psi_n, P_n)\}$$

  states that θ, a formula encoding one or more actions and some fluent preconditions, has the effect of causing exactly one of the fluent conjunctions $\psi_1, \ldots, \psi_n$ with probabilities $P_1, \ldots, P_n$ respectively.
- the *s-proposition*

  $$\theta\ \textbf{senses}\ F\ \textbf{with-accuracies}\ M$$

  states that θ has the effect of sensing fluent F with an accuracy given by the confusion matrix M.
- the *o-proposition*

  $$A\ \textbf{occurs-at}\ I\ \textbf{with-prob}\ P\ \textbf{if-holds}\ \theta$$

  states that action A is known to be occurring at instant I with probability P, but only if its preconditions encoded in θ are satisfied.
- the *p-proposition*

  $$A\ \textbf{performed-at}\ I\ \textbf{if-believes}\ (\theta, \bar{P})$$

  states that action A is performed by the agent at instant I if its state of belief in formula θ at instant I falls in the (open, half-open or closed) interval $\bar{P}$.

A *domain description* in EPEC is a collection of these propositions satisfying some integrity constraints (e.g. exactly one i-proposition must belong to any domain description). We are not going to describe these constraints formally here, but the interested reader can find them in [29], [30].

EPEC has a possible-worlds semantics where each world represents a possible evolution of the world from the initial state and is weighted according to the propositions in the domain descriptions. Four implementations of EPEC are available and/or under active development, and can answer queries regarding what is true (and with what probability) in a given domain. Two of them are optimised for the non-epistemic fragments of EPEC, called PEC+. While the exact implementation (written in clingo [31]) exhaustively works out all the possible worlds and their associated weights, the approximate implementation (written in the probabilistic programming language Anglican [32]) samples a user-defined number of worlds using Anglican's built-in Markov Chain Monte Carlo sampling capabilities, and uses the obtained sample to approximate the probability of a query. Similarly, there are two implementations of EPEC (including the epistemic fragment) to deal with exact and approximate inference.

2) Knowledge Base: The following simple domain description demonstrates some features of EPEC:

$$\textit{Engagement}\ \textbf{takes-values}\ \langle \textit{false}, \textit{true} \rangle \tag{1}$$

$$\textbf{initially-one-of}\ \{(\{\textit{Engagement}\}, 1)\} \tag{2}$$

$$\textit{EEG}\ \textbf{senses}\ \textit{Engagement}\ \textbf{with-accuracies} \begin{pmatrix} 0.7 & 0.3 \\ 0.4 & 0.6 \end{pmatrix} \tag{3}$$

$$\textit{Watching}\ \textbf{senses}\ \textit{Engagement}\ \textbf{with-accuracies} \begin{pmatrix} 0.8 & 0.2 \\ 0.1 & 0.9 \end{pmatrix} \tag{4}$$

$$\textit{Cutscene}\ \textbf{causes-one-of}\ \{(\{\textit{Engagement}\}, 0.9), (\emptyset, 0.1)\} \tag{5}$$

$$\forall I, \textit{EEG}\ \textbf{performed-at}\ I \tag{6}$$

$$\forall I, \textit{Watching}\ \textbf{performed-at}\ I\ \textbf{if-believes}\ (\textit{Engagement}, [0, 0.7]) \tag{7}$$

$$\forall I, \textit{Cutscene}\ \textbf{performed-at}\ I\ \textbf{if-believes}\ (\textit{Engagement}, [0, 0.5]) \tag{8}$$

These propositions aim at describing an automated system used to detect the degree of engagement of the patient, which can sound a dedicated alarm to raise the patient's level of engagement if this falls below an appropriate threshold. In this example, *Engagement* is a boolean valued fluent which at every instant can take values *true* or *false* (proposition (1)) and initially the patient is known to be fully engaged (proposition (2)). Propositions (3) and (4) specify the confusion matrices associated with the actions of *EEG* and *Watching* through the system's sensors, while proposition (5) defines what the expected effects of playing the *Cutscene* is, i.e. raising the patient's level of engagement in 90% of the cases. *EEG* is continually performed (proposition (6)), whereas *Watching* is only performed if belief in *Engagement* falls below the 0.7 threshold (proposition (7)). Finally, the *Cutscene* is only played if *Engagement* is believed to have fallen below the 0.5 threshold (proposition (8)).

The AI engine therefore established if this is more appropriate to select more exercises or if it is necessary to apply an attention recovery strategy. Indeed, also the modality of execution of an exercise (e.g., its speed) can be adjusted with respect to the children profile.

In this simple case, *EEG* and *Watching* are thought to be independent. In EPEC it is also possible to model dependency between actions. For instance, consider a case in which a high-res camera is also employed to perform engagement detection, and consider its associated action *HiResWatching*. This could be modelled by appropriately reworking proposition (4)'s precondition and adding the two propositions:

$$\mathit{HiResWatching} \wedge \neg \mathit{Watching}\ \textbf{senses}\ \mathit{Engagement} \quad \textbf{with-accuracies} \begin{pmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{pmatrix} \tag{10}$$

$$\mathit{HiResWatching} \wedge \mathit{Watching} \quad \textbf{with-accuracies} \begin{pmatrix} 0.91 & 0.09 \\ 0.07 & 0.93 \end{pmatrix} \tag{11}$$

Notice that, although *HiResWatching* is more accurate than *Watching* (compare propositions (4) and (10)'s matrices) these two actions are correlated, and the confusion matrix in proposition (11) reflects this.

IV. CONCLUSIONS

We have presented the system architecture designed for the AVATEA project to manage adaptive rehabilitation exercises and provide therapists with interpretable feedback about the session. Also, we have presented the hybrid approach to combine the use of explicit rules with probabilistic management of noisy data sources, like automated classifiers working on streamed sensor data. The system will autonomously manage rehabilitation exercises and will react to social feedback coming from young users during a gamified experience. After the end of the session, the system will provide a detailed report about the session to support therapists in evaluating children improvement and design further interventions.

ACKNOWLEDGEMENTS

This work has been partially supported by MIUR within the POR Campania FESR 2014-2020 AVATEA "Advanced Virtual Adaptive Technologies e-hEAlth" research project.

REFERENCES

[1] P. Wilson, D. Green, K. Caeyenberghs, B. Steenbergen, and J. Duckworth, "Integrating new technologies into the treatment of cp and dcd," *Current Developmental Disorders Reports*, vol. 3, no. 2, pp. 138–151, Jun 2016.

[2] N. Hocine, A. Gouaïch, and S. A. Cerri, "Dynamic difficulty adaptation in serious games for motor rehabilitation," in *Games for Training, Education, Health and Sports*, S. Göbel and J. Wiemeyer, Eds. Cham: Springer International Publishing, 2014, pp. 115–128.

[3] M. S. Cameirao, S. Bermudez i Badia, E. D. Oller, and P. F. M. J. Verschure, "Using a multi-task adaptive vr system for upper limb rehabilitation in the acute phase of stroke," in *Virtual Rehabilitation*, Aug 2008, pp. 2–7.

[4] S. L. Oviatt and P. Cohen, "Perceptual user interfaces: multimodal interfaces that process what comes naturally," *Communications of the ACM, Volume 43, Issue 3*, pp. 45–53, 2000.

[5] M. Pirovano, R. Mainetti, G. Baud-Bovy, P. L. Lanzi, and N. A. Borghese, "Self-adaptive games for rehabilitation at home," in *2012 IEEE Conference on Computational Intelligence and Games (CIG)*, Sep. 2012, pp. 179–186.

[6] R. M. Ryan and E. L. Deci, "Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being." *American psychologist*, vol. 55, no. 1, p. 68, 2000.

[7] T. Belpaeme, P. E. Baxter, R. Read, R. Wood, H. Cuayáhuitl, B. Kiefer, S. Racioppa, I. Kruijff-Korbayová, G. Athanasopoulos, V. Enescu *et al.*, "Multimodal child-robot interaction: Building social bonds," *Journal of Human-Robot Interaction*, vol. 1, no. 2, pp. 33–53, 2012.

[8] N. Vernadakis, M. Papastergiou, E. Zetou, and P. Antoniou, "The impact of an exergame-based intervention on children's fundamental motor skills," *Computers & Education*, vol. 83, pp. 90–102, 2015.

[9] Y. Wu, Y. Cai, and D. Tu, "A computerized adaptive testing advancing the measurement of subjective well-being," *Journal of Pacific Rim Psychology*, vol. 13, 2019.

[10] A. Origlia, P. Cosi, A. Rodà, and C. Zmarich, "A dialogue-based software architecture for gamified discrimination tests." in *GHITALY@ CHItaly*, 2017.

[11] M. Mulcahey, S. M. Haley, T. Duffy, P. Ni, and R. R. Betz, "Measuring physical functioning in children with spinal impairments with computerized adaptive testing," *Journal of pediatric orthopedics*, vol. 28, no. 3, p. 330, 2008.

[12] C.-Y. Huang, L.-C. Tung, Y.-T. Chou, H.-M. Wu, K.-L. Chen, and C.-L. Hsieh, "Development of a computerized adaptive test of children's gross motor skills," *Archives of physical medicine and rehabilitation*, vol. 99, no. 3, pp. 512–520, 2018.

[13] A. Skarlatidis, A. Artikis, J. Filippou, and G. Paliouras, "A Probabilistic Logic Programming Event Calculus," *Theory and Practice of Logic Programming*, vol. 15, pp. 213–245, 3 2015.

[14] A. Skarlatidis, G. Paliouras, A. Artikis, and G. A. Vouros, "Probabilistic Event Calculus for Event Recognition," *ACM Transactions on Computational Logic (TOCL)*, vol. 16, no. 2, p. 11, 2015.

[15] V. Belle and H. J. Levesque, "Reasoning about Discrete and Continuous Noisy Sensors and Effectors in Dynamical Systems," *Artif. Intell.*, vol. 262, pp. 189–221, 2018.

[16] M. Van Lambalgen and F. Hamm, *The proper treatment of events*. John Wiley & Sons, 2008, vol. 6.

[17] P. M. Nadkarni, L. Ohno-Machado, and W. W. Chapman, "Natural language processing: an introduction," *Journal of the American Medical Informatics Association*, vol. 18, no. 5, pp. 544–551, 2011.

[18] T. Sato, "A statistical learning method for logic programs with distribution semantics," in *IN PROCEEDINGS OF THE 12TH INTERNATIONAL CONFERENCE ON LOGIC PROGRAMMING (ICLP95*. Citeseer, 1995.

[19] F. Riguzzi, *Foundations of Probabilistic Logic Programming*, ser. River Publishers Series in Software Engineering. River Publishers, 2018.

[20] M. Richardson and P. Domingos, "Markov logic networks," *Machine Learning*, vol. 62, no. 1, pp. 107–136, Feb 2006.

[21] L. De Raedt, A. Kimmig, and H. Toivonen, "Problog: A probabilistic prolog and its application in link discovery." in *IJCAI*, vol. 7, 2007, pp. 2462–2467.

[22] L. Nigay and J. Coutaz, "A design space for multimodal systems: concurrent processing and data fusion," in *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. ACM, 1993, pp. 172–178.

[23] M. Turk, “Multimodal interaction: A review,” *Pattern Recognition Letters*, vol. 36, pp. 189–195, 2014.
[24] S. Rossi, E. Leone, M. Fiore, A. Finzi, and F. Cutugno, “An extensible architecture for robust multimodal human-robot communication,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov 2013, pp. 2208–2213.
[25] Z. Cao, G. Hidalgo, T. Simon, S. Wei, and Y. Sheikh, “Openpose: Realtime multi-person 2d pose estimation using part affinity fields,” *CoRR*, vol. abs/1812.08008, 2018.
[26] N. Jaques, O. O. Rudovic, S. Taylor, A. Sano, and R. Picard, “Predicting tomorrows mood, health, and stress level using personalized multitask learning and domain adaptation,” in *Proceedings of IJCAI 2017 Workshop on Artificial Intelligence in Affective Computing*, ser. Proceedings of Machine Learning Research, N. Lawrence and M. Reid, Eds., vol. 66, 20 Aug 2017, pp. 17–33.
[27] S. Rossi, A. Caso, and F. Barile, *Trends in Practical Applications of Agents, Multi-Agent Systems and Sustainability: The PAAMS Collection*. Cham: Springer International Publishing, 2015, ch. Combining Users and Items Rankings for Group Decision Support, pp. 151–158.
[28] S. Rossi, L. Bove, S. Di Martino, and G. Ercolano, “A two-step framework for novelty detection in activities of daily living,” in *Social Robotics*. Cham: Springer International Publishing, 2018, pp. 329–339.
[29] F. A. D’Asaro, A. Bikakis, L. Dickens, and R. Miller, “Foundations for a Probabilistic Event Calculus,” in *LPNMR*, 2017, pp. 57–63.
[30] F. A. D’Asaro, “Probabilistic Epistemic Reasoning About Actions,” Ph.D. dissertation, University College London, February 2019.
[31] M. Gebser, R. Kaminski, B. Kaufmann, and T. Schaub, “Clingo = asp + control: Preliminary report,” vol. arXiv:1405.3694v1, 2014.
[32] D. Tolpin, J. W. van de Meent, H. Yang, and F. Wood, “Design and Implementation of Probabilistic Programming Language Anglican,” *arXiv preprint arXiv:1608.05263*, 2016.