HHAI 2024: Hybrid Human AI Systems for the Social Good
F. Lorig et al. (Eds.)

doi:10.3233/FAIA240204

A Conversational Agent for Structured Diary Construction Enabling Monitoring of Functioning & Well-Being

Piek VOSSEN [a,1], Selene BÁEZ SANTAMARÍA [a] and Thomas BAIER [a]
[a] *Vrije Universiteit Amsterdam*
ORCiD ID: Piek Vossen https://orcid.org/0000-0002-6238-5941, Selene Báez Santamaría https://orcid.org/0000-0003-3808-2614, Thomas Baier https://orcid.org/0000-0002-1260-3373

Abstract. We describe a Hybrid Intelligence agent that constructs a personal diary through conversation. The diary is represented in an episodic Knowledge Graph as a timeline of events, where the communication is driven by the *information need* of the agent. We argue that such a structured diary provides valuable information to contextualize physical, social and mental functioning and well-being for medical research and monitoring. We provide details on the formal model and implementation and demonstrate the communication by our first baseline agent. Our code is available under the MIT license on GitHub: `https://github.com/leolani/cltl-diary-parent`.

Keywords. conversational agent, timeline reconstruction, health, monitoring

1. Introduction

Knowing about somebody's well-being and functioning implies knowing about their life. Hybrid Intelligence solutions that service individual needs and goals could benefit from keeping track of personal circumstances and developments, as well as people's future plans. Especially in the context of medical support systems, in which agents need to monitor people on a regular basis, knowledge about someone's life is crucial for understanding their perspective and making personal decisions. In many cases, such monitoring takes place over an extended period to measure the impact of, e.g. treatment, training or coaching and to be able to detect trends or unexpected patterns that can alert caretakers. There are various benefits from automating such monitoring: it can take place continuously or on a regular basis to obtain more data points, it is less invasive, and it offers more privacy as people do not need to communicate personal information and feelings to other people directly but only abstract or generalized data can be shared.

In a medical context, automated monitoring is mainly done using sensors to detect Activities of Daily Life (ADL) [1,2,3]. Such monitoring is, however, limited to activities

[1]Corresponding Author: Piek Vossen, `p.t.j.m.vossen@vu.nl`

that can be detected and it lacks a deeper interpretation and perspective of the patient.[2] We therefore propose a Hybrid Intelligence approach in which a communicative agent talks with patients on a regular basis to construct a structured diary: similar to what a friend would do. As Figure 1 illustrates, our agent uses an ontology and a reconstructed timeline to ask questions regarding a user's personal events.

Such an agent can 1) register activities (both physical and social), including those that are not observable by technological devices, 2) get more details about these activities by asking follow-up questions, 3) get information on the patient's perspective on these activities, and 4) get information on activities that did not happen or are planned in the near future. By having a regular friendly conversation, an agent can create a timeline for somebody's past life and future plans. In addition to reconstructing the timeline, the agent can also directly ask for specific conditions that need to be monitored for medical purposes. In comparison to a self-written diary, an agent-driven structured diary has a more explicit and standardized representation that can better generalize over the lives of different people and can stimulate and remind people regularly to provide more information, that might be relevant but missing, on their life and details of each event.

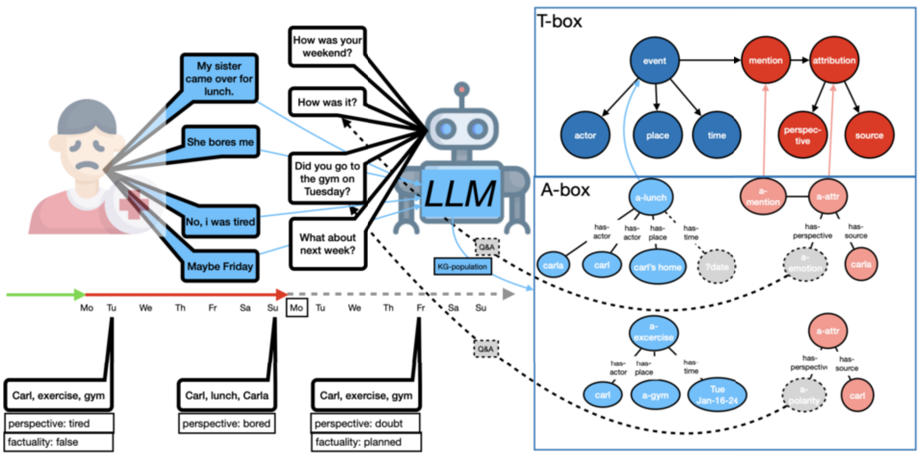

Figure 1. Structured diary reconstruction by a conversational agent that uses a fine-tuned LLM to understand the patient input as a populated graph (A-box) given a predefined ontology (T-box). Gaps in knowledge drive communication to get more information and perspectives and to validate information. Dotted grey circles represent unknown event properties that can be addressed through questions. The timeline consists of three periods: the period before the previous conversation (green), the period since the previous conversation (red) and the future (dotted). These periods define potential gaps to be filled with events.

In this paper, we describe a conversational agent for timeline reconstruction implemented in the Leolani framework [4,5] using its capability to pro-actively drive conversations, similar to [6] who also use an ontology to drive communication but for understanding cooking recipes. We begin by describing three medical Hybrid-Intelligence use-cases in Section 2, arguing that structured diaries provide valuable contexts for the monitoring of functioning and well-being of patients and being able to generalize over patients. We proceed to define the problem in Section 3 and explain how we formalize and implement the problem in Section 4.

[2]In this paper, we use patient synonymous to user and fragile elders at home

2. Background and motivation

Various studies address the use of sensor technology to detect activities in the context of eHealth systems. In order to define what these activities are and to facilitate sharing of data, Woznowski et al. [1] describe a hierarchical ontology for activities of daily living (ADL). Their ontology provides a good starting point for physically perceivable activities in homes. Majumder et al [2] provide a survey of research on smart home healthcare technologies. They specifically address the use of sensors and actuators that facilitate remote monitoring of the home environment (such as temperature, humidity, and smoke in the home) as well as important physiological signs (such as heart rate, body temperature, blood pressure and blood oxygen level), and activities of the occupants. Oladinrin et al [3] investigate the perception of professionals on the use of smart home technology to improve and enhance the "ageing-in-place" of elderly residents. They conclude that the development of smart home technology for health support is complex and there are various mismatching perspective. It is essential to involve the target patients and elderly both in the design but also in the monitoring as such.

Sensor-based monitoring is limited to activities that can be detected and it lacks a deeper interpretation and perspective of the patient. Activities that cannot be observed or picked up by sensor, either because they are out of reach (e.g. going out for shopping or a walk) or because they represent socially defined activities (e.g. working or planning a holiday), will be missed. Furthermore, the patient is not actively involve to provide further details, complement (contextualizing) or correct detected activities, and to express their appreciation or perspective on the activity. Finally, the patient may have certain expectations and plan with respect to the future that is not yet observable but may be important for understanding their perspective.

Deeper and more comprehensive information from the patient can be obtained by questionnaires, as shown by the InterRAI initiative.[3] They develop questionnaires as internationally standardised instruments to obtain medical data from both caretakers and patients. Similarly, patient data acquisition and monitoring is applied within protocols for Trauma Care Pathways [7] and revolving door processes in healthcare [8]). Such protocols are however restricted in the amount of data points than can be acquired over time. Especially obtaining long-term data for patients is a challenge as it requires collecting personal information beyond the scope of the direct medical care. Furthermore, questionnaires are not tailored to personal lives and are tedious to fill in, which requires discipline from the patients. Diaries and questionnaires are one-way communication processes that do not allow further interaction to clarify, correct or elaborate.

Conversational agents may offer a number of advantages to sensor-based technology and questionnaires, especially with respect to follow-up questions, addressing errors, inconsistencies, conflicts and uncertainties and possibly other information related to medical conditions. They can communicate regularly and conveniently, producing more data points in time by actively reminding and stimulating patients. The interaction can be made more engaging and empathetic than writing diaries and filling questionnaires and more robust and standardized compared to open and free diaries kept by patients.

In the next subsections, we describe three use cases under investigation as a further motivation in which structured diaries can provide context for lifestyle monitoring, clinical treatment and research purposes.

[3] https://interrai.org

2.1. Patient recovery and functioning patterns

A 2019 UN report indicates that over 46% of individuals aged 60 years and above suffer with disabilities, thereby impinging on their ability to independently perform activities of daily life. As a corollary of the surge in chronic conditions, frailty, and disabilities among older adults, healthcare utilization has also witnessed a sharp increase, as these individuals are frequently (re)hospitalized due to their disabilities [9]. Multimorbidity, i.e., the presence of more than one chronic condition in one person, is affecting 65% of individuals aged 65-84 years and approximately 82% of those aged 85 years and above. Frail older adults with multimorbidity are prone to unplanned hospitalization by e.g., disease exacerbation, disability and falls. Hospitalizations, in turn, can further exacerbate morbidity and disability with subsequent loss of self-dependency.

Counteracting this vicious circle requires close monitoring on multiple domains of health and functioning outside hospital conditions in order to obtain insight in critical conditions in their way living. However, this requires much effort from patients and an already scarcely available workforce. A conversational agent can support this monitoring. By obtaining sufficiently rich data over time, the agent can provide more context to professionals judging the medical and mental conditions of people but also offers the opportunity to infer the functioning level directly from the patient as defined by the World Health Organization in the ICF standard: the International Classification of Functioning, Disability and Health (ICF). The agent could leverage existing medical Large Language Models (LLM) fine-tuned for ICF level classification [10,11]. Likewise, the agent can keep track of the functioning level on different ICF categories such as capability to *walk*, *eat*, *concentrate*, their *energy* or *mood* while conversing about their lives. Any changes in functioning over time can be signaled and produce an alert for caretakers.

2.2. Diabetes

Type 2 Diabetes Mellitus is another more specific growing health concern, characterized by chronic hyperglycemia resulting from insulin resistance and impaired insulin secretion. While it is relatively well understood that lifestyle changes are effective [12], adherence to treatment remains a challenge. Reasons for this are the chronic nature of the required behavioural changes, the time investment required, and the fact that different types of change need to be maintained simultaneously [13]. Effective lifestyle recommendations should therefore incorporate the patient's context, preferences and values in order to maximize adherence [14].

An AI system can support patients in life style changes by monitoring their ADLs and by providing coaching. Such an AI system could also support in GP consults by proposing potential lifestyle change that are likely to be adopted by the patient and become a long lasting change. The AI system would provide this advice based on the patient recent and long term diary and thus account for their drives, barriers, current lifestyle and past attempted changes. Similar to the previous use case, the same structured diary can be used as a context for specific monitoring of conditions relevant for diabetes.

2.3. Toxicity

Cancer therapy is often associated with toxicity, especially for chemotherapies that regrettably affect dividing cells outside the tumor. For brain cancer glioblastoma, toxicity is commonly seen after application of temozolomide, a chemotherapeutic that is part of the standard therapy. In around 13% of the patients this leads to toxicities because of blood dysfunction caused by an gradual depletion of thrombocytes (i.e. blood platelets) during therapy. Patients that are treated for cancer over a longer period of time, can suffer from varies adverse side effects of the treatment. These effects show up in their daily life after the actual treatment itself. Due to the complexity of the medication, it is very difficult to obtain sufficient data on these effects. Monitoring patients well-being and perception of life can provide valuable data at scale to evaluate the impact, especially when this is paired with the clinical treatment according to a very strict schedule. This will shed light on additional care (unplanned hospitalization) and adverse events or changes in functioning that are not directly related to the treatment.

The proposed agent can be guided by the strict predefined treatment plan and activities to provide a primary calendar which is the starting point for inquiring about any (medical) events and conditions outside the protocol. Furthermore, the agent can inquire about the perspective of the patient on these events, within and outside the protocol. This information is crucial for medical professional to learn about the impact of treatments but also to adjust treatment to personal conditions and perceptions.

3. Problem description

We frame our structured diary goal as a timeline reconstruction problem that is driven by the need of the agent to become knowledgeable about both the past and any future plans. We define three types of unknowns: 1) possible and probable events and their properties, 2) possible and probable perspectives on these events and 3) factuality of events (realis or confirmed, irrealis or not confirmed and explicitly denied [15]). Note that in our approach there is no independent way of knowing what really happened except from the direct source of information provided. In other words, what the patient confirms or claims is true by lack of other evidence and what the patient denies did not happen. In order to obtain a more advanced interpretation, the agent needs to consult other sources or make independent observations.

Figure 1 shows an example of our agent interpreting events and their properties as well as their perspective in terms of associated emotions, certainty, beliefs, expectations and denials. We define a timeline as a temporal container in which events are to be placed relative to the time of the conversation. Different periods are distinguished within this container: 1) before the previous conversation, 2) the gap between the current and previous conversation and 3) the future beyond the current conversation.

The problem is defined as to use conversation to become informed about events and their properties and put them on a timeline in corresponding periods until a certain level of density and saturation is reached. In order to address this problem, the agent needs to reason over the populated Knowledge Graph and timeline to estimate missing knowledge, find knowledge that needs certification, or project to the future as expected knowledge. Missing knowledge is defined in the T-box and A-box ontology as what should be the case, what is likely the case given previous instances (analogy, probabilities), or what is expected in the temporal containers (e.g. based on habits and hobbies). The agent needs to ask questions based on the need to reconstruct and validate the timeline, given a definition what is sufficient and reasonably obtainable from the targeted person. Therefore, the agent needs to model what is sufficient and obtainable as stopping criteria.

4. Model and Implementation

In this section we introduce how we modelled and implemented our solution to this problem, focusing on the intents that drive the agent.

4.1. Model

For modeling our problem, we use the following concepts:

- Events (E): occurrences of an activity grounded in time.
- Activity (Y), Actors (A), Time (T) and Place (P) as properties of events.
- Mentions (M): expressions that make reference to events, actors, place and time.
- Claims (C): statements made by a source.
- Sources (S): the interlocutors or a third party making claims.
- Perspectives (V): epistemic beliefs, sentiments, or emotions of the source towards a claim.
- Timeline (L): a period on a temporal ruler, used to ground events in time.
- Now (N): the current encounter of an agent and a human.
- Previous (R): the previous encounter of an agent and a human.
- History (H): the period on the timeline before the previous encounter (R).
- Gap (G): the period in between the previous encounter (R) and the current encounter (N).
- Future (F): the period after the current encounter (N).
- Density (D): proportion of events per period given a threshold of the expected events .
- Saturation (U): proportion of knowledge obtained for an event given the expected properties.

Representing events We use the Simple Event Model [16] as a formal model for defining events in somebody's life. Every instance of an `sem:Event` $e_i \in (E)$ is defined by a tuple (Y_e, A_e, P_e, T_e), which captures the *what*, *who*, *where* and *when* through the relations `sem:hasActor`, `sem:hasPlace`, and `sem:hasTime`.

Representing activities For ADL, we rely on the WHO's International Classification of Functioning, Disability and Health (ICF[4]). ICF measures health and disability at both individual and population levels. In this work we use the category of *Activities and participation*, which is subdivided into 9 sub-chapters such as *Domestic life*: *Preparing meals*, *Doing housework*, or *Community, social and civic life*: *Play*, *Sports*, *Socializing*. The ADL events are included in the T-box as subclasses of `n2mu:Activity`[5] while individual events are `rdf:instanceOf` of `sem:Event` and a specific ADL type. We can represent constraints on the actors, place, and time using ontology types such as people, animals, places, artifacts, materials, dates, periods, e.g. washing is done by `n2mu:Person`, using `n2mu:Water` and `n2mu:Soap` in places of type `n2mu:Bathroom` or `n2mu:Kitchen`.

Representing conversations We use the GRaSP framework [17] to represent mentions of events in conversation. Mentions are tokens (words) within utterances from a source. A mention $m_i \in (M)$ will have a `grasp:denotes` relation with a `claim` $c_j \in (C)$. Each mention has also an attribution `grasp:hasAttribution` that specifies the source and their perspective on a claim, where the same claim can be made multiple times by the same or different sources representing different perspectives. As perspective values, we currently use `grasp:polarity` (categorical: realis, irrealis and denied), `grasp:certainty` (scalar, be-

[4] https://www.who.int/standards/classifications/international-classification-of-functioning-disability-and-health

[5] The namespace *n2mu* stands for Nice-To-Meet-You and was developed in [4]

tween 0 and 1) `grasp:sentiment` (scalar between -1 and 1) and `grasp:emotion` (categorical, 27 emotions defined in [18]). The lunch example in Figure 1 is modelled as follows:

Event representation as RDF triples

```
grasp:claim1{
:lunch_01 rdf:instanceOf sem:Event, icf:Lunch .
:lunch_01 sem:hasActor n2mu:Carl, n2mu:Carla .
:lunch_01 sem:hasPlace n2mu:Carl_home .
:lunch_01 sem:hasTime :21_01-2024 .
}
```

```
grasp:mention1 grasp:denotes grasp:claim1 .
grasp:mention1 grasp:hasAttribution n2mu:attr1 .
n2mu:attr1 prov:hasSource n2mu:Carl .
n2mu:attr1 grasp:Emotion :DISAPPROVAL .
n2mu:attr1 grasp:Certainty .8 .
n2mu:attr1 grasp:Polarity :Realis .
```

What is not mentioned in the conversation and therefore not included in this representation is, for example, what food and drinks they had for lunch. If this is defined in the ontology as possible actors of the activity, it can be used to drive questions. The density threshold $d_i \in (D)$ is the stop criterion to ask for more events and the saturation threshold $u_i \in (U)$ defines the upper bound for properties (SEM relations and GRaSP perspectives) to be learned about each event. Further prioritization for asking questions is based on learned probabilities from the history (H).

4.2. Implementation

Following [4,5], our agent tracks interaction using an event-bus as temporally-grounded data through which incoming signals are stored sequentially as so-called topics. Any processing module can be connected to the event-bus as a service that is sensitive to certain input topics and pushes its interpretation as output of a new topic, which can then become the input for another module. A module becomes active whenever the conditioned input topics appear on the event-bus. The architecture provides the flexibility to build any input-output pipeline using the defined modules by defining and connecting topics.

Conversation takes place when modules interpret speech or text input signals from the user (patient) and produce output topics as system responses rendered as speech or text output signals. System responses in a simple chat system can be generated directly, as in an Eliza [19] setup or a generative LLM that is prompted, or through more complex processing using multiple modules. In our case, we use an episodic Knowledge Graph (eKG) [4] to store the interpretation of the user input because 1) we want to reason more precisely about and have control over the interpretation of the signal given a person's life and the current conversation and 2) the reasoning defines the information need or hunger of the agent and drives the communication. Besides, an eKG does not require user-specific fine-tuning with training data and can be combined with LLMs to generate responses with lower risk of hallucination.

Using a Believe-Desire-Intent protocol [20], the platform utilizes modules that define global high-level desires. This module calls other (generic) modules that interpret incoming signals as beliefs represented in the eKG and produce intents to achieve these goals.

4.3. Conversational intent

For the agent under consideration, we define the following desires and intents:

1. **high level desires** driven by the timeline:
   Did what was expected/planned happen: *Did you had that lunch with your sister?*

What happened since the last conversation: *How was your weekend?*
What is planned to happen: *What are your plans for tomorrow?*

2. **mid level intents** driven by necessary/possible/probable properties in the ontology:
   Who, what, when, where: *What did you had for lunch?*
   Analogies: *Did she stay long again?*
   Probabilities: *Did you drink wine?*
3. **low level intents** driven by possible and likely perspectives:
   emotion: *How was it for you?*
   certainty: *Are you sure?*
   conflicts: *But you told me before that...*

The overall flow of the interaction is initiated by determining the gap (G) between the now (N) and the previous encounter (R). During previous encounters, expectations were projected into (G) that should be addressed first, after which the agent can inquire about other potential and habitual events. If (G) is filled according to density and saturation conditions $((D), (U))$, the agent will inquire about the expectation and plans for the future using SPARQL queries to reconstruct the timeline. For each event, the known properties and perspectives are extracted from the eKG using further queries. Figure 4.3 gives a schematic example of multiple mentions of the same event with different source perspectives (*Irrealis* and *Denied*), and a mention of as future event on a timeline.

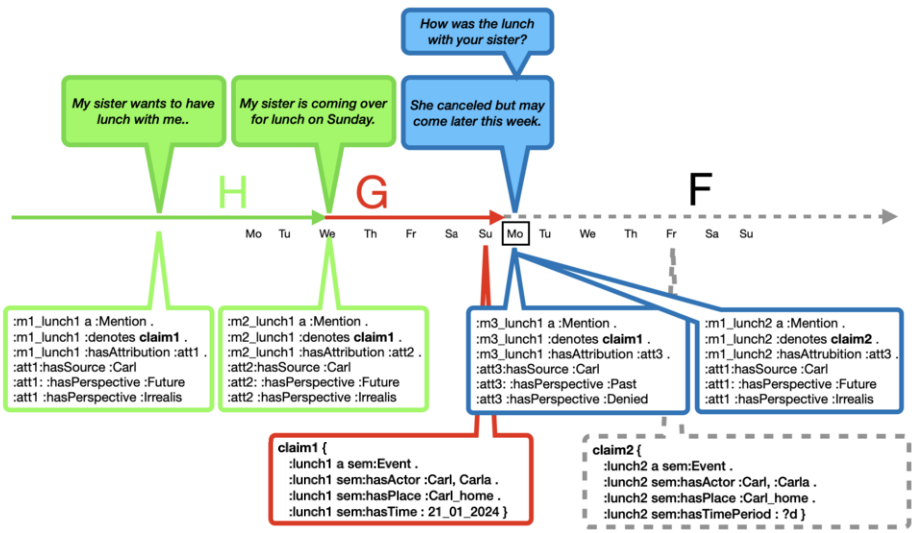

Figure 2. Multiple mentions of the same event with different perspectives and a mention of a future event on a timeline. The event `lunch1`, represented as a set of triples `claim1`, is mentioned multiple times in different conversations: `m1_lunch1, m2_lunch1, m3_lunch1`. The event itself is grounded in time on Sunday, 21:01:2024 through `m2_lunch1`, which lays in time period (G). In the current conversation, the patient denies the event from happening which changes the perspective from *Irrealis* (not known as a fact) to *Denied*. At the same time, the patient introduces a new event in (F) with an undefined date.

Addressing events that may have happened but are not mentioned yet and therefore not introduced on the timeline, is done through open questions for a period (What did you do lately?, How was your weekend?) and by iterating further (What else did you do?) or on the basis of habitual patterns from the history (Did you go to the gym?).

These questions do not require further SPARQL queries but are driven by patterns and expectations from the past and applying density and saturation conditions only.

For each ADL event considered in the high-level processing, the agent starts a sub-loop to ask for necessary and expected relations (who, what, where, when or mid-level intents), and another sub-loop to ask about epistemic and emotional perspectives relations (low-level intents). For generating questions and statements from the agent, we use the conversational history in combination with the relevant structured data from the eKG to prompt a generative LLM (decoder). For processing responses and answers to questions from the patient, we use encoder LLMs fine-tuned for emotion detection, dialogue act classification, triple extraction.

Density and Saturation criteria Events and their properties can be necessary, possible and probable. Likewise, we define the density and saturation thresholds within reasonable limits but also using the personal history. Density and saturation can be reached at all three levels of Desire and Intents, where priorities can be defined in different ways. The current implementation uses configurable thresholds for the required density of events per period and for the number of properties per event. We further use configurable thresholds to give-up after a number of attempts.

New information that is integrated into the eKG typically generates a plethora of graph patterns to respond to. From these, the agent needs to select the most effective ones to reach the density and saturation criteria. Success depends on the collaboration with the user as well as on the conversational skills of the agent (natural language understanding and generation, dialogue management). Implementation-wise, selection can be random, top-down combined with random, scripted or based on policies learned through reinforcement learning. We currently implemented a baseline with scripted interaction, proceeding as described above. Within the high-level script, a specific event is targeted randomly but this can be adapted to relevance, habitual patterns or personal preferences.

5. Conclusions and future work

In this paper, we described a conversational agent that constructs a structured diary of a person's life and perspective. Such a structured diary provides rich and valuable contextual information for medical monitoring, coaching and research. The current baseline of the agent uses a script over three levels of desires and intents. In future work, we will further enhance the agent using reinforcement learning with learned policies for intent selection and carry out empirical testing of different components and policies. For training the system, we will generate synthetic conversational data prompting generative models with synthetic structured event data as shown here. For testing, we plan to use human - human conversations.

6. Acknowledgements

This research was funded by the Vrije Universiteit Amsterdam and the Netherlands Organisation for Scientific Research (NWO) via the Spinoza grant (SPI 63-260) awarded to Piek Vossen and the Hybrid Intelligence Centre via the Zwaartekracht grant (024.004.022).

References

[1] Woznowski P, Burrows A, Diethe T, Fafoutis X, Hall J, Hannuna S, et al. SPHERE: A sensor platform for healthcare in a residential environment. Designing, Developing, and Facilitating Smart Cities: Urban Design to IoT Solutions. 2017:315-33.

[2] Majumder S, Aghayi E, Noferesti M, Memarzadeh-Tehran H, Mondal T, Pang Z, et al. Smart homes for elderly healthcare—Recent advances and research challenges. Sensors. 2017;17(11):2496.

[3] Oladinrin OT, Mesthrige JW, Ojo LD, Alencastro J, Rana M. Smart Home Technologies to Facilitate Ageing-in-Place: Professionals Perception. Sustainability. 2023;15(8):6542.

[4] Baez Santamaria S, Baier T, Kim T, Krause L, Kruijt J, Vossen P. EMISSOR: A platform for capturing multimodal interactions as Episodic Memories and Interpretations with Situated Scenario-based Ontological References. In: Donatelli L, Krishnaswamy N, Lai K, Pustejovsky J, editors. Proceedings of the 1st Workshop on Multimodal Semantic Representations (MMSR). Groningen, Netherlands (Online): Association for Computational Linguistics; 2021. p. 56-77. Available from: https://aclanthology.org/2021.mmsr-1.6.

[5] Baier T, Santamaria SB, Vossen P. A modular architecture for creating multimodal agents. arXiv preprint arXiv:220600636. 2022.

[6] Steels L, Verheyen L, van Trijp R. An experiment in measuring understanding. In: HHAI2022: Augmenting Human Intellect. IOS Press; 2022. p. 241-2.

[7] Bourke JB, Dobrovolny J, Eaton M, Ferrante T, Smith M. Complex trauma care pathway: results of a 12-month pilot. The Permanente Journal. 2021;25.

[8] Williamson AE, Mullen K, Wilson P. Understanding "revolving door" patients in general practice: a qualitative study. BMC family practice. 2014;15:1-11.

[9] Stoeldraijer L, Van Duin C, Huisman C. Bevolkingsprognose 2017–2060: 18, 4 miljoen inwoners in 2060. Centraal Bureau voor de Statistiek. 2017;2017:1-19.

[10] Meskers CG, van der Veen S, Kim J, Meskers CJ, Smit QT, Verkijk S, et al. Automated recognition of functioning, activity and participation in COVID-19 from electronic patient records by natural language processing: a proof-of-concept. Annals of Medicine. 2022;54(1):235-43.

[11] Kim J, Verkijk S, Geleijn E, van der Leeden M, Meskers C, Meskers C, et al. Modeling Dutch Medical Texts for Detecting Functional Categories and Levels of COVID-19 Patients. In: Proceedings of the Thirteenth Language Resources and Evaluation Conference; 2022. p. 4577-85.

[12] Brown A, McArdle P, Taplin J, Unwin D, Unwin J, Deakin T, et al. Dietary strategies for remission of type 2 diabetes: a narrative review. Journal of Human Nutrition and Dietetics. 2022;35(1):165-78.

[13] Galaviz KI, Narayan KV, Lobelo F, Weber MB. Lifestyle and the prevention of type 2 diabetes: a status report. American journal of lifestyle medicine. 2018;12(1):4-20.

[14] Koenigsberg MR, Bartlett D, Cramer JS. Facilitating treatment adherence with lifestyle changes in diabetes. American family physician. 2004;69(2):309-16.

[15] Saurí R, Pustejovsky J. FactBank: a corpus annotated with event factuality. Language resources and evaluation. 2009;43:227-68.

[16] Van Hage WR, Malaisé V, Segers R, Hollink L, Schreiber G. Design and use of the Simple Event Model (SEM). Journal of Web Semantics. 2011;9(2):128-36.

[17] Fokkens A, Vossen P, Rospocher M, Hoekstra R, van Hage W. Grasp: Grounded representation and source perspective. In: Proceedings of KnowRSH, RANLP-2017 workshop, Varna, Bulgaria; 2017. .

[18] Demszky D, Movshovitz-Attias D, Ko J, Cowen A, Nemade G, Ravi S. GoEmotions: A Dataset of Fine-Grained Emotions. In: 58th Annual Meeting of the Association for Computational Linguistics (ACL); 2020. .

[19] Weizenbaum J. ELIZA—a computer program for the study of natural language communication between man and machine. Communications of the ACM. 1966;9(1):36-45. Available from: https://dl.acm.org/doi/10.1145/365153.365168.

[20] Bratman M. Intention, plans, and practical reason. 1987.