Multi-Agent System for Recruiting Patients for Clinical Trials

Samhar Mahmoud
King's College London
London, UK
samhar.mahmoud@kcl.ac.uk

Gareth Tyson
Queen Mary
London, UK
gareth.tyson@qmul.ac.uk

Simon Miles
King's College London
London, UK
simon.miles@kcl.ac.uk

Adel Taweel
King's College London
London, UK
adel.taweel@kcl.ac.uk

Tjeerd Vanstaa
London School of Hygiene &
Tropical Medicine
London, UK
tjeerd.vanstaa@lshtm.ac.uk

Michael Luck
King's College London
London, UK,
michael.luck@kcl.ac.uk

Brendan Delaney
King's College London
London, UK
brendan.delaney@kcl.ac.uk

ABSTRACT

Clinical trials are widely adopted for the purpose of evaluating medical research. In particular, they are used to study various aspects of medical science, as well as being a vital stage in the deployment of new drug treatments. However, a review of the UK Medical Research Council found that only 31% of trials actually recruited to their planned target, with 30–40% of costs arising during the recruitment phase alone. This is mainly due to the challenges that are involved in designing the clinical trial, and recruiting the required number of patients for this trial within a certain time-frame. Both tasks create significant overhead as they are slow and costly. In response, we propose a multi-agent architecture that helps ease the process of recruiting patients for clinical trials. This paper presents a results from a deployment of the architecture, showing that it succeeds in recruiting a sufficient number of patients for multiple clinical trials. The results also show that recruitment is better for some trials than for others, due to the differing trial requirements and recruitment processes.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Design, Performance

Keywords

Healthcare Application, Multi-Agent, Evaluation

1. INTRODUCTION

Clinical trials are the gold standard by which medical research is evaluated. They involve the controlled testing of treatments on patients who match certain criteria, e.g. age, gender, ailment. The stringent nature of these criteria, however, mean that many trials are unsuccessful in recruiting sufficient patients. A review of the UK Medical Research Council found that only 31% of trials actually recruited to their planned target, with 30–40% of costs arising during the recruitment phase alone [7], as discovering and contacting eligible potential recruits are logistically and legally challenging. Consequently, many research projects take far longer to complete than is desirable (or, in the worst case, not at all).

Currently, patient recruitment is performed in a laborious manner, which is ill-suited for situations where many specialised patients are required. It often involves a human recruitment agent visiting clinics in an attempt to locate suitable patients (e.g. asking practitioners or searching local medical records). This creates significant overhead as it is both slow and costly, as well as non-scalable for most trials.

With this challenge in mind, we began an effort to build a system capable of facilitating clinical researchers in more effectively recruiting patients for their trials. Through an extensive requirements elicitation process, in collaboration with the UK's National Health Service, we discovered that the recruitment process is not only complex, but also crosses the boundaries of many different legal entities. Typical trials can involve several organisations, each containing many different professionals, including members of universities, government bodies, companies, regulation frameworks, hospitals and various other primary care units. Each organisation has different aims, as well as a variety of internal protocols that must be followed. We believe this type of environment suits, ideally, the strengths of agent-based design and, as such, this paper details the design, implementation and deployment of an agent-based recruitment system for clinical trials, called *ePCRN-IDEA*.

The principle behind our work is simple: to replace human recruitment agents with autonomous software agents. We have modelled the processes of clinical trial recruitment using the GAIA methodology [15], designing an agent-based system consisting of all entities required to successfully discover, recruit and monitor patients. We have implemented and deployed the majority of our full design, placing our software agents in 31 primary care clinics in the UK. Put simply, these agents monitor the patients seen

within each clinic in real time, attempting to identify those who could be recruited onto active trials. Whenever an eligible patients visits a clinic, the General Practitioner (GP) is notified and asked to undertake a simple online procedure to recruit the individual. In collaboration with a number of other agents distributed in various organisations, these agents have already recruited 230 patients. This paper details our experiences during this process, describing both our design and prototype implementation.

The rest of the paper is structured as follows. The background to the research is discussed in Section 2, before introducing our model of clinical trials in Section 3. Section 4 details the ePCRN-IDEA recruitment system, Section 5 presents a prototype implementation, and Section 6 contains an evaluation. We summarise the outcomes in Section 7.

2. BACKGROUND

We begin by exploring the area of clinical trial recruitment before discussing, more generally, agents in healthcare. Following this, we provide a brief description of the GAIA agent development methodology, which we chose as the most appropriate for designing ePCRN-IDEA.

2.1 Clinical Trial Recruitment

Clinical trials are a challenging stage in the research of clinicians due to the complexity of recruiting patients for participation. Many types of trials can suffer from such difficulties; for instance, trials that have potential recruits who are widely distributed over many clinics are extremely difficult to recruit for due to the intensive resource requirements. Studies show that 30% of participating clinics fail to even recruit a single patient [10]. This can be exacerbated by the complex eligibility requirements, or trials that require immediate actions, e.g. a change of drug treatments.

Clinical trial recruitment is performed by first defining *eligibility criteria* that stipulate the exact characteristics that make a patient eligible for participation, e.g. gender, age, ailments etc. It is then necessary to discover patients who match the criteria, before contacting and recruiting them. Traditionally, locating such patients is achieved using one or more of the following approaches:

- *Advertising and public relations*: This involves using posters, adverts and brochures to advertise eligibility criteria directly to practitioners and patients.

- *Recruiters*: This involves sending human recruiters to clinics, usually after feasibility modelling, analysis and site selections, in an attempt to discover patients who match the eligibility criteria.

- *Practitioners*: This involves doctors meeting periodically (often GPs) to discuss patient treatments and potential trials in an attempt to spot eligible patients during consultation.

These methods are time consuming and expensive, particularly for trials that have high patient targets, complex eligibility criteria, rare diseases or involve emergency cases. This has led to the development of Clinical Trial Alert (CTA) systems, which automatically alert practitioners to the eligibility of a patient when they are in consultation. Such systems allow the practitioner to immediately discuss the trial with the patient, to enable instant recruitment in a trusted environment. This process is achieved by automatically comparing patient information against computable eligibility criteria in real-time during consultations. However, most of these systems are bespoke [4, 1, 2], recruiting patients to a single trial within a single clinic. Other similar techniques have also seen only limited large-scale testing [13]. The challenge of designing generic systems that can handle multiple trials is exacerbated by the complexity of interactions between the various organisations. Trust and security issues, for example, are significant, often hampering the ability to deploy non-bespoke systems. Similarly, technical challenges such as scalability becomes notable when expanding deployments across multiple trials and clinics (eligibility criteria can be extremely complicated to compute). All of these have meant that CTA systems still remain in their infancy.

2.2 Agent Based Healthcare Systems

The use of agent-based systems for healthcare has seen widespread investigation. Nealon et al. [9] discusses 11 areas in which agent technology is being applied to improve healthcare in Europe, including integration of heterogeneous patient records [8], control of cardiac pacing, and monitoring the elderly using agent-based tele-assistance. For example, MAID [3] is an agent-based system for integrating heterogeneous data sources within a hospital environment. The studied hospital had 24 departments, each using their own information systems. To address this, agents were constructed to interoperate with each system to monitor changes and retrieve data for insertion into a central repository. In a subsequent work, HealthAgents [5] also enabled decision support, specifically for diagnosing brain tumours.

Agent-based systems have also been proposed for handling distributed expertise, such as using agents to enable better communication between healthcare workers based on ambient information, e.g. their role, location etc. [12], as well as using agents to remotely monitor patients [6, 11]. These systems often involved data analysis. S(MA)^{2}D, for instance, uses statistical analysis to cluster patients into similar groups [11]. This ability to scalably perform data analysis in real time also shows potential for enabling the type of eligible patient identification discussed previously. Despite this, there has been little study into using agents to improve clinical trial recruitment.

2.3 The GAIA Methodology

In our work, we have applied the GAIA agent-oriented software engineering methodology to develop our agent-based system. The GAIA methodology [15] was proposed to guide the process of developing a multi-agent system from analysis to design. For brevity, we focus here on one particular model used in the methodology: the *role model*. A role describes a particular agent behaviour, and is defined in the form of a schema, which consists of the following parts: *description* is used to briefly express the purpose of the role; *protocols* and *activities* are the actions and tasks that the role are provided with to achieve its goal and to communicate with other roles; *permissions* specify the information that the role has access to or able to produce; and, *responsibilities* define the generalised behaviour pattern of the role (liveness) as well as states of affair that the role must maintain or bring to existence (safety). Instances of roles are shown as part of our design below.

3. MODELLING CLINICAL TRIALS

The first problem in designing a healthcare system like ePCRN-IDEA is the diversity of languages and standards used within the different key organisations. This problem has plagued many, often unsuccessful, healthcare systems. To address this, our work started with the construction of a formal model capable of capturing the key attributes of clinical trials. This was chosen as it is the only body of information that must consistently be understood across all organisations. The key design goal of the model has been to allow easy translation from existing data structures, better enabling
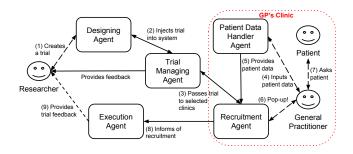
Figure 1: Recruitment System Architecture

independent organisations to interoperate [14]. In contrast, the data format of individual bi-lateral information flows can be defined by the agents involved.

A full specification of the model is beyond the scope of this paper, however, here we provide a brief overview to assist in general understanding. The model can be separated into a number of sub-models that collectively capture all important aspects of a trial's design. These can be summarised as follows:

- *Trial Description:* This model holds a description of a trial, e.g. the title, overview, funding body etc. It allows practitioners to be presented with details of the trial during the recruitment phase. Much of the data in this model is therefore free-text, which can be used to easily generate descriptions for various entities (e.g. GPs, public websites, researchers).

- *Eligibility Criteria:* This model represents a formal computable set of criteria for patient eligibility. It defines the characteristics required of any participating patients. The model is SQL-like in nature, allowing complex sets of predicates to be executed over any fields in a patient's medical record. These predicates can be used to both include and exclude a patient from recruitment. Typical predicates would include age, gender and ailments (there are existing terminology standards for describing ailments). Often, patients with particular existing ailments or treatments would also be excluded. Note that this model does not stipulate *when* or *where* the criteria should be executed to compute eligible patients. For example, the criteria could be used in real-time as patient data is entered or, alternatively, in batches every month from a centralised database.

- *Recruitment Model:* This model stores information about the recruitment process itself. It stipulates details such as how many patients should ideally be recruited, which clinics are authorised, which GPs are authorised, how patients should be recruited if they are interested, the priority of recruiting for the particular trial and recruitment timeplans. This is an important component in allowing the agents to reason over the finer-grained needs of the trial's recruitment process.

These models were defined through collaboration between several parties, including doctors, developers and medical and computer science researchers. The models have been captured within XML schemas that (as will be later discussed) are shared amongst all agents and organisations, allowing each one to map their own data to the trial schema.

4. SYSTEM DESIGN

The previous section has provided an overview of the way ePCRN-IDEA models clinical trials. We exploit this model to enable agents in different organisations to interoperate. This section now details the design of ePCRN-IDEA, which we have decomposed into a set of agents responsible for a variety of tasks. During our design process, we have followed the GAIA development methodology; as such, we also include, where appropriate, formal GAIA descriptions of the agents. Due to space constraints, the interaction between all the different agents is depicted on the overall system design diagram and not through the interaction model diagrams of GAIA.

Note that, although, we have developed our prototype for the UK, we intend our design to be applicable for many different healthcare systems. As such, in this section, we wish to strictly separate our design from implementation and, thus, endeavour to avoid specific implementation details. Instead, details of our prototype realisation of this design are delayed until Section 5.

4.1 Overview

As discussed in Section 1, the basic idea behind ePCRN-IDEA is to place software agents within the primary care clinics (more accurately, on GPs' computers). These agents interact with a number of other humans and software agents to discover, recruit and monitor patients who are eligible for any trials that exist within the system. Figure 1 illustrates the architecture of the overall system, where arrows indicate the interactions between agents, i.e. the flow of information. To provide greater insight, the links are also annotated with example operations that occur through these interactions. Note that these are just examples and do not provide an exhaustive set of interactions.

The process of trial recruitment begins with the *Designing Agent*. This is an agent that resides within the organisation of the *researcher*. The researcher is a human operator responsible for defining new trials. The Designing Agent works with the researcher to generate an instantiation of the clinical trial model described in Section 3. The Designing Agent then passes the newly created trial model to one or more *Trial Managing Agents*. Many of these exist throughout the healthcare system available in the country; typically, these would be distributed in a geographical sense. In the UK, for example, healthcare is split between multiple geographical management groups termed primary care trusts (each would have one agent). The Trial Managing Agents are responsible for coordinating in which clinics recruitment should be performed (for each trial). Once they have decided this, they pass the new trial model to a set of *Recruitment Agents*. These exist on every GP's computer in every clinic. The Recruitment Agent monitors the data entry by the GP to decide if any patient matches the eligibility criteria for the trials it is aware of. If a match is observed, the Recruitment Agent generates a graphical pop-up asking the GP to recruit the patient. Obviously, this requires real-time access to patient information, which is provided by a *Patient Data Handler Agent*. This also sits on the GP's computer, but is owned and managed by the company that controls the patient's medical records.[1] Once someone has been recruited, an *Execution Agent* is contacted to initiate the trial with the patient. In the simplest case, this might involve retrieving information (e.g. a blood pressure reading), whereas in other cases it could involve complex interventions (e.g. drug treatments) accompanied by proactive data collection.

---

[1]In the UK, several commercial vendors provide Electronic Health Record databases. Clinics are largely free to select their preferred one.

Below, we describe these agents in more detail, explaining their responsibilities and relation with each other.

4.2 Designing Agent

Before any recruitment can take place, it is necessary for the researcher to provide a comprehensive description of the trial for the system. Most notably, this includes the formal eligibility criteria for patient inclusion. This process is managed by the Designing Agent, which we detail in Table 1 using a GAIA role model.

The Designing Agent resides within the organisation of the researcher, e.g. a university or pharmaceutical company. It acts as the representative of the researcher, and brokers interactions between the researcher and the wider ePCRN-IDEA system. Its key task is to receive the trial description as identified by the researcher (*ReceiveCriteria*) and express it in a systematic manner (*DesignSpecification*). In order to confirm that the specification expresses the researcher's generic criteria, a series of communications must take place between the researcher and Designing Agent until the researcher's approval is received (*AcknowledgeSpecification*).

Following this, the Designing Agent must submit the newly created trial to the rest of the system (*DistributeSpecification*). This will often involve a variety of negotiations; the most prominent example relates to the monetary charges that are issued (many entities expect payment for their role, e.g. clinics). This negotiation can also involve adaptation of the trial description within the ranges stipulated by the researcher, e.g. lowering the desired number of patient recruits. Note that often the researcher would want to be involved in authorising these negotiations.

4.3 Trial Managing Agent

Once a Designing Agent has created a new trial description, it is required to submit it to the system. This submission, and subsequent negotiation, is received by one or more Trial Managing Agents (*ReceiveTrial*). The functionality of this agent is captured in the GAIA role model, shown in Table 2. The Trial Managing Agents are responsible for coordinating the placement of trial recruitment across all clinics under their control (*DistributeTrial,*).

The Trial Managing Agents exist throughout the heathcare system(s) resident in the country. In the UK, there is a single integrated healthcare system (the National Health Service). This, however, is separated into several organisational units termed primary care trusts; each has independent management, legal constraints and protocols to follow. To address this diversity, a Trial Managing Agent exists in each. Their key responsibility is to select in which clinics a trial should be recruited for. Due to both legal and scalability reasons, it is impossible for every trial to exist in every clinic. As such, this must be a well thought out decision: some clinics may have the potential to recruit tens of patients, whereas others will be unable to recruit even a single one. Similarly, complex inter-dependencies between trials must also be managed (*RepriortiseTrials* and *SuspendTrials*). For example, trials with extremely similar eligibility criteria should not recruit in the same clinic; this is because the one with the higher priority would likely prevent the other from recruiting (thereby making its presence redundant). These decisions therefore require interaction between the different Trial Managing Agents, as well as with the clinics themselves.

Following the deployment of the trials in the clinics, the Trial Managing Agent is responsible for monitoring recruitment in their area (*MonitorRecruitment*). This results in a feedback loop, allowing trials to be re-allocated to different clinics where appropriate. The Trial Managing Agents must also report back to the Designing Agent so that the researcher can follow progress (*ProvideFeedback*). This can also result in adaptation of the trial description

itself. For example, one of the pilot trials in ePCRN-IDEA had its criteria adapted after several weeks due to problems with recruitment.

4.4 Recruitment Agent

Once the Trial Managing Agents have decided which clinics and GPs should perform recruitment for a given trial, the details must be sent to the clinics. This is received by the Recruitment Agents (*ReceiveTrial*), which reside on the computers of any GPs authorised for recruitment. The functionality of the Recruitment Agent are captured in the GAIA Role model presented in Table 3.

The Recruitment Agents act as mediators between the patient, the GP and the trial itself. Its key responsibility is to discover and recruit patients in the clinic for the trials it knows of. Whenever a patient visits a clinic, the GP enters information into their patient database, e.g. about new diagnoses. This information is passed to the local Recruitment Agent (*ReceivePateintData*) (via the Patient Data Handler Agent) and used to compute the patient's eligibility for any known trials (*CheckEligibility*). For example, a drug trial looking to test a new treatment on people with joint pain would clearly need patients who are complaining of joint pain.

Once a Recruitment Agent has computed the trials that a patient is eligible for, it is responsible for ensuring that a recruitment takes place. Through user interviews and experimentation, we have found that GPs are often very reluctant to expend time and cognition on dealing with recruitment whilst a patient is in consultation. As such, the Recruitment Agent is driven towards *only* presenting trials that it believes will result in recruitment. Whenever a set of eligible trials is computed, the Recruitment Agent is tasked with selecting which to display, i.e. prioritising them (*RankTrials*). This is based on trial-specific factors (e.g. the trial's urgency), as well as GP-specific factors (e.g. the types of trials the GP is interested in).

If the patient is interested in recruitment, the GP is can confirm it with the Recruitment Agent (*RecruitPatinet*). Following this, the Recruitment Agent must contact the Trial Managing Agent to inform it of the update (*NotifyManagingAgent*), and the Execution Agent to start the follow up process (*NotifyExecutionAgent*).

4.5 Patient Data Handler Agent

To enable the Recruitment Agent to fulfil its task, it is necessary to provide it with real-time information about patients during their consultation. Although simplistic in many domains, this is extremely difficult in healthcare due to the sensitivity of data. The responsibility of managing patient data is therefore given to the Patient Data Handler Agent. The functionality of this agent is presented in the GAIA role model, shown in Table 4.

The Patient Data Handler Agent resides alongside every Recruitment Agent in the clinic. It is owned and managed by the Electronic Health Record database proprietor, as a separate entity to the other key organisations (note that several proprietors exist in the UK). Whenever requested (*ReceiveRequest*), the Patient Data Handler Agent provides information about the patient in question (*ExtractData* and *InformRecruitment*). This requester is most frequently the local Recruitment Agent itself, however, it may also be a Trial Designing Agent or a Trial Managing Agent, which also use patient data. The key task undertaken by the Patient Data Handler Agent is ensuring privacy is not undermined. This is achieved by executing various privacy policies on data requests, restricting access to authorised users. Im

4.6 Execution Agent

The previously described agents have been primarily responsible for recruiting patients into a clinical trial. However, the process

Table 1: The Designing Agent Role Model

Role Schema: Designing Role
Description: This role involves designing clinical trial specifications that satisfy with researcher's requirements.
Protocols and Activities: ReceiveCriteria, DesignEligibilityCriteria, DesignRecruitmentCriteria, DesignSpecification, DistributeSpecification, AcknowledgeSpecification
Permissions: reads *researcherCriteria* // the criteria of the trial generates *trialSpecification* // the complete design of trial
Responsibilities: Liveness: DesingTrial= (ReceiveCriteria.DesignSpecification.DistributeSpecification)$^{\omega}$ DesignSpecification = (DesignEligibilityCriteria.DesignRecruitmentCriteria.AcknowledgeSpecification)$^{\omega}$ Safety: equivalent(ClinicalTrial , researcherCriteria)

Table 2: The Trial Managing Role Model

Role Schema: Trial Managing Agent
Description: This role involves monitoring the recruitment for clinical trials within certain geographical location.
Protocols and Activities: ReceiveTrial, MonitorRecruitment, DistributeTrial, ReprioretiseTrials, SuspendTrials, CheckStatus, ProvideFeedback
Permissions: reads *trialRecruitmentStatus* // the current recruitment status of a trial deactivate *trialID* // the identifier of clinical trial
Responsibilities: Liveness: Trial = (ReceiveTrial.DistributeTrial.MonitorRecruitment$^{\omega}$.ProvideFeedback)$^{\omega}$ MonitorRecruitment = (CheckStatus.(ReprioretiseTrials|SuspendTrials).InformRecruitmentAgent)$^{\omega}$ Safety: ReachedTrialDeadline = true => trialRecruitment >= targetRecruitment

does not end there. Depending on the trial, there may be multiple phases that the patient needs to go through: few trials are only about collecting some static information about the patient. Consequently, it is necessary to monitor the execution of the trial across its multiple phases, making sure that all the phases' goals are met and the result are collected and analysed. The Execution Agent is therefore responsible for executing the trial itself; a formal description of its functionalities is provided in Table 5. A huge number of Execution Agents could exist (in theory, one per trial). In the simplest case, they could request the retrieval of periodic information about a given patient. However, this could also involve instructing GPs to perform various (changing) interventions (*CheckPhases*). Clearly, the final responsibility of the Execution Agent is collating the collected data (*ObtainResult*) and providing it to the researcher (*InformResearcher*).

5. PROTOTYPE IMPLEMENTATION

The previous section has presented the design of ePCRN-IDEA. We now detail our prototype realisation of this design, which we have deployed across 31 clinics in the UK. So far, we have implemented a subset of the overall architecture, replacing certain agents with simpler alternatives. Several organisations have been involved in this process, including universities, clinics, companies and government organisations. Each, as discussed in the previous section, plays an integral role in achieving the overall system goal. The agents have been implemented as follows:

- *Designing Agent*: The Designing Agent has been implemented within a Workbench toolkit that is provided to researchers. This allows the authoring and submission of new trials. Currently, the Designing Agent only exists in a single research organisation: the UK's Clinical Practice Research Datalink (CPRD). They are currently running a number of active clinical trials that ePCRN-IDEA is facilitating with. The negotiation process between the Designing Agent and the rest of ePCRN-IDEA is currently quite rudimentary, with all trials that match authorisation policies being accepted. Similarly,

all monetary transactions are performed external to the system. Clearly, the prioritised future work for this agent is enhancing its negotiation capabilities to reduce the need for any external interactions between parties.

- *Trial Managing Agent*: The Trial Managing Agent has been implemented in Java, with support from a MySQL back-end database. This is used to conveniently exchange trial information in a secure way. Due to the legal constraints on such systems, these types of security considerations were vital for achieving deployment. Due to the relatively small number of pilot clinics (31), only one Trial Managing Agent currently has been instantiated. It operates within King's College London (technical coordinators of the system). Within the prototype, its current responsibilities centre on the distribution of trials, as well as the collection of progress information. The latter is used to activate and deactivate recruitment, as well as feed information back to the Designing Agent for interested researchers. The key future work of interest to us is scaling the number of Trial Managing Agents up, and enabling their interaction, which currently we do not support. Unfortunately, due to the need to first increase the number of clinics, we are investigating these principles via simulation.

- *Recruitment Agent*: The Recruitment Agent has been implemented in Java and has been used by 124 GPs so far. Despite obvious deployment challenges, it has largely been well received with many practitioners keen to be involved. This agent has been the key focus of our work so far, and currently supports all the tasks detailed in the earlier design. Figure 2 provides a screenshot of the Recruitment Agent's graphical pop-up. Through experimentation, we discovered that most GPs were only prepared to be presented with a single trial alert (rather than the original list of several that we provided). As such, the Recruitment Agent only displays the highest priority one. The GP is then allowed to register several responses, as shown in Figure 2. All inter-agent interaction is then automatically handled, without requiring any

Table 3: The Recruitment Role Model

Role Schema: Recruitment Agent
Description: This role involves checking the eligibility of patients for clinical trial.
Protocols and Activities: ReceiveTrial, ReceivePateintData, CheckEligibility, RankTrials, NotifyGP, RecruitPatinet, NotifyManagingAgent, NotifyExecutionAgent
Permissions: reads: *trialData* // The trial criteria *patientReleventData* // the data of the patient generate: *notification* // notification about patient's eligibility
Responsibilities: Liveness: PateintRecruitment = (ReceivePateintData.CheckEligibility.RankTrials.RecruitPatinet.NotifyManagingAgent.NotifyExecutionAgent)$^{\omega}$ Safety: RecruitedPatinet => EligiblePatient

Table 4: The Patient Data Handler Role

Role Schema: Patients Data Handler Role
Description: This role involves maintaining patient data across multiple consultations, and extracting trial relevant data and send it to the recruitment agent
Protocols and Activities: ReceiveRequest, ExtractData, InformRecruitment
Permissions: reads: *patientData*
Responsibilities: Liveness: PateintDetection = (ReceiveRequest,CheckPatient.InformRecruitment)$^{\omega}$

external actions by the user.

- *Patient Data Handler Agent*: The Patient Data Handler Agent has so far been implemented within a single Electronic Health Record system called Vision. This has been performed by the proprietor, a company called INPS. Vision covers about 20% of all UK clinics. So far, it does not perform the advanced functionality in our design, instead, passing a fixed set of patient attributes to the Recruitment Agent whenever a new patient record is opened or changed during a consultation. To pass information from the Patient Data Handler Agent to the Designing Agent and Trial Managing Agent, a central database hosted at CPRD is used as an intermediary. Once again, this is necessary to satisfy the tight security requirements on our deployment: for a prototype it was considered unacceptable to send confidential patient information to multiple agents in different organisations in an entirely automated way (note that CPRD has a trusted position, hence this choice). Clearly, our key future work is therefore to overcome these regulatory challenges and enable more direct inter-agent communication of such data. Further, we wish to expand the privacy controls within the Patient Data Handler Agent as a key component of enabling this.

- *Execution Agent*: The Execution Agent is an agent that is developed on a per-trial basis. This is because different trials can have wildly different needs for their execution. We currently have the functionality of the Execution Agent split between a bespoke website developed by a third party company and a special department within CPRD. Collectively, these are responsible for performing the management tasked detailed in the design. For example, they study the status of the patients, and issue appropriate instructions to GPs for any necessary interventions (e.g. changes in drug treatments).

6. EVALUATION

The previous section has described our prototype deployment of ePCRN-IDEA. Due to the novelty of such real-world deployments, we choose to focus our evaluation on the results gained

Table 6: Description of Deployed Trials

Trial Name	Description
eLung	The eLung trial targets patients aged over 40 with a medical history of Chronic Obstructive Pulmonary Disease (COPD) who, in the opinion of the GP, had an acute exacerbation of COPD with an increase of non-purulent sputum volume, who did not require immediate referral to specialist care for treatment of COPD exacerbation and consented to participation.
RETROPRP	Retropro is a pragmatic point-of-care trial, and target patients with age over 40, a 20% or greater 10-year risk of developing CVD, primary hypercholesterolemia and consent to participation.
FLU-CATs	The FLU-CATs study evaluates and refines community assessment tools for use as decision aids in preparation for any future influenza pandemic or similar event where health-seeking behaviour exceeds clinical capacity.

through this. More specifically, we ask the question: *has ePCRN-IDEA been successful in improving recruitment?* Due to space constraints, we do not explore the more technical components (e.g. overheads).

So far, the Recruitment Agent has been installed in 31 clinics with 124 GPs involved. We have been successful in initiating three (diverse) clinical trials within the system, summarised in Table 6. We have found that *all* GPs involved recruited at least one patient. Overall, ePCRN-IDEA has generated 3204 alerts, with 230 patients being recruited (7.1%). Although, at first, this may not seem significant, it is actually far above the recuitment levels one would likely achieve through traditional means, e.g. advertisements. Interestingly, we also found that 45 of these alerts resulted in the GP deciding that the patient was not actually eligible. This verifies the point raised by a number of clinicians throughout the development of ePCRN-IDEA, which is that the agents must *always* oper-

Table 5: The Execution Role Model

Role Schema: Execution Agent
Description: This role involves monitoring the execution of a clinical trial after a patient has been recruited.
Protocols and Activities: ReceiveTrial, ReceivePatient, CheckPhases, ObtainResult, InformResearcher
Permissions: reads: trialData
Responsibilities:
Liveness: TrialExecution = (ReceivePatient.CheckPhases .ObtainResult.InformResearcher)$^\omega$
Safety: PhaseDeadlineReached= true => PhasesResultObtained = true

ate within the confines of human supervision. In terms of patients' responses (other than acceptance) 30 patients were not interested and 14 patients wanted time to think. This means that the majority of alerts were actually ignored by the GP, suggesting that many unwanted interrupts could (and should) have been avoided by the Recruitment Agent. We believe this warrants further work to ensure that Recruitment Agents are able to learn *when* to best interrupt the GP with an alert (if at all). That said, despite being offered, none of the participating GPs requested that the recruiting agent be stopped. A summary of results is shown in Table 7.
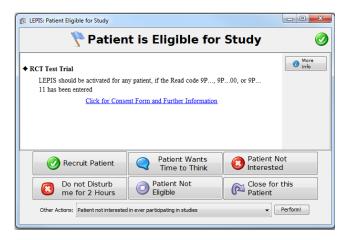

Figure 2: GP Notification Respond Interface

Whereas overall we noted positive results from the trials, we also observed significant variations between the behaviour of the different trials. To explore this, Figure 3 shows the number of notifications generated per-month between January and September of 2013 (note that *FLU-CATS* only started in the second month). It can be seen that ePCRN-IDEA was effective at generating notifications for all trials, particularly *RETROPRO* during the later months. This dramatic increase occurred due to a re-definition of the eligibility criteria. We found that ePCRN-IDEA's management mechanisms (e.g. the Trial Manager Agent) were vital for achieving this re-definition, which was brought about because many GPs were not entering the properly formatted information into their Electronic Health Record database, instead using quick free-text (*RETROPRO* was the only trial that needed this). To address the issue, the trial researchers adapted the eligibility criteria to use a pre-computed list of 'at-risk' patients instead. This therefore generated pop-ups whenever the patient's record was opened. As can be seen, this resulted in a spike in recruitment, achieving over 40 new recruits at its peak. This highlights effectively the flexibility of the trial models used, as well as the ability of the architecture to monitor and react to behaviour in the system (in collaboration with its human operators). The relationship between the notifications and recruitment can also be seen in Figure 4. As can be observed from this figure,
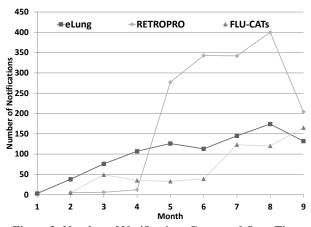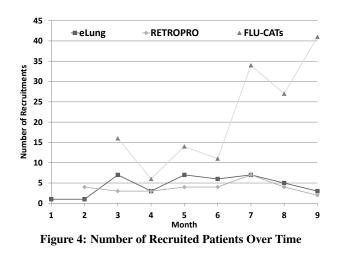

Figure 3: Number of Notifications Generated Over Time

Figure 4: Number of Recruited Patients Over Time

some trials managed to recruit more patients than others, although overall a roughly positive trend can be seen across all cases. The trial with the greatest number of recruits is *FLU-CATs*, as this has only a very low overhead for the patients.

Through these results, we would argue that there have been very real benefits offered by ePCRN-IDEA. By decomposing the system into self contained agents, this has also been achieved in a way that has satisfied the various stakeholders by empowering them to (internally) operate in a way that matches their personal needs and norms. Practically speaking, without this, the deployment would likely have not taken place.

7. CONCLUSION AND FUTURE WORK

This paper has detailed the design, implementation and deployment of an agent-based system called ePCRN-IDEA. The research started in an attempt to devise novel ways in which technology

Table 7: Detailed Observation of the three trials

Trial ID	eLung	RETROPRP	FLU-CATs
Start date	21-01-2013	19-01-2013	28-01-2013
Result reporting date	30-09-2013	30-09-2013	30-09-2013
Number of flags	914	1589	569
Number of patients recruited	40	31	149
Number of patients not interested	10	6	14
Number of patients not eligible	25	18	2
Number of patients who need time to think	1	6	7

could facilitate the recruitment of patients to clinical trials. We quickly realised, however, that the complexity and cross-organisational nature of the domain would make traditional software architectures difficult to use in practice. As such, we have built and deployed a comprehensive agent-based design that has already managed to recruit 230 patients to three pilot trials in operation. Our evaluation so far has centred on the ability of ePCRN-IDEA to enable superior recruitment. It has verified that both the technology and the underlying principles are correct, and that systems such as this have a great potential. It has also highlighted the flexibility of our architecture, allowing the agents and human operators to adapt to reflect new observations (e.g. adapting eligibility criteria in response to user behaviour).

A critical point to raise, however, is the disparity between our system design and the prototype implementation. Whilst we have built a range of sophisticated algorithms into our agent design, we have not been able to integrate all of them into our deployment. The reasons for this have been diverse, and warrant discussion in themselves. At one end of the scale, in some cases, we did not have sufficient manpower to realise their complexity. However, in the majority of cases, this exclusion was the product of legal and security constraints brought about by the nature of the domain. The focus for future work is therefore to further our prototype by replacing the simpler agent implementations with that originally detailed in the design. Of most priority are the Patient Data Handler Agent and the Execution Agent. Beyond this, we are also continually working on all the existing agent prototypes to increase their sophistication; of most interest is the deployment of multiple Trial Management Agents to better realise their potential for collaboration when managing the placement of trial recruitment amongst the different trials. Clearly, alongside this technical future work, we are also working towards the expansion of our pilot by including increasing numbers of clinics and trials. Our final aim is then to achieve a full system evaluation that explores the intricacies of the agents in in wild.

8. REFERENCES

[1] L. B. Afrin, J. C. Oates, C. K. Boyd, and M. S. Daniels. Leveraging of open emr architecture for clinical trial accrual. *AMIA Annual Symposium proceedings AMIA Symposium AMIA Symposium*, 2003:16–20.

[2] A. J. Butte, D. A. Weinstein, and I. S. Kohane. Enrolling patients into clinical trials faster using realtime recuiting. *Proceedings of the AMIA Symposium*, pages 111–115, 2000.

[3] R. Cruz-Correia, P. Vieira-Marques, P. Costa, A. Ferreira, E. Oliveira-Palhares, F. Araújo, and A. Costa-Pereira. Integration of hospital data using agent technologies - a case study. *AI Commun.*, 18, August 2005.

[4] P. J. Embi, A. Jain, J. Clark, S. Bizjack, R. Hornung, and C. M. Harris. Effect of a clinical trial alert system on physician participation in trial recruitment. *Archives of Internal Medicine*, 165(19):2272–2277, 2005.

[5] H. González-Vélez, M. Mier, M. Julià-Sapé, T. Arvanitis, J. García-Gómez, M. Robles, P. Lewis, S. Dasmahapatra, D. Dupplaw, A. Peet, C. Arús, B. Celda, S. Van Huffel, and M. Lluch-Ariet. Healthagents: distributed multi-agent brain tumor diagnosis and prognosis. *Applied Intelligence*, 30, 2009.

[6] V. Koutkias, I. Chouvarda, and N. Maglaveras. A multiagent system enhancing home-care health services for chronic disease management. *Information Technology in Biomedicine, IEEE Transactions on*, 9(4):528–537, dec. 2005.

[7] A. McDonald, R. Knight, M. Campbell, V. Entwistle, A. Grant, J. Cook, D. Elbourne, D. Francis, J. Garcia, I. Roberts, and C. Snowdon. What influences recruitment to randomised controlled trials? a review of trials funded by two uk funding agencies. *Trials*, 7(1):9, 2006.

[8] M. Nagy and M. Vargas-Vera. *Towards an Automatic Semantic Data Integration: Multi-agent Framework Approach*. Chapter in Sematic Web. In-Tech Education and Publishing KG, 2010.

[9] J. Nealon. Agents applied in health care. *AI Communications*, page 22, 2005.

[10] R. Nitkin. Patient recruitment strategies. Training workshop conducted by National Institutes of Health, Bethesda, Md, 2003.

[11] A. Rammal, S. Trouilhet, N. Singer, and J.-M. Pécatte. An adaptive system for home monitoring using a multiagent classification of patterns. *Int. J. Telemedicine Appl.*, 2008:3:1–3:8, January 2008.

[12] M. D. Rodríguez, J. Favela, A. Preciado, and A. Vizcaíno. Agent-based ambient intelligence for healthcare. *AI Commun.*, 18:201–216, August 2005.

[13] B. L. Rollman, G. S. Fischer, F. Zhu, and B. H. Belnap. Comparison of electronic physician prompts versus waitroom case-finding on clinical trial enrollment. *Journal of General Internal Medicine*, 23(4), 2008.

[14] G. Tyson, A. Taweel, S. Zschaler, T. V. Staa, and B. Delaney. A model-driven approach to interoperability and integration in systems of systems. In *Proc. of Workshop on Model-Based Software and Data Integration (MBSDI)*, 2011.

[15] M. Wooldridge, N. R. Jennings, and D. Kinny. The gaia methodology for agent-oriented analysis and design. *Journal of Autonomous Agents and Multi-Agent Systems*, 3:285–312, 2000.