Dorneich, Michael C., Whitlow, Stephen D., Olson, Eric G., Anhalt, David, and Monteith, Tracy (2010). “The Combat Causal Reasoner Approach to Robotic Control”, *Proceedings of the Human Factors and Ergonomics Society Conference*, San Francisco, CA. September 27-October 1.

The Combat Causal Reasoner Approach to Robotic Control

Michael C. Dorneich, Stephen D. Whitlow
Honeywell Laboratories
Eric Olson, David Anhalt, Tracy Monteith
SAIC

This paper describes an approach to autonomous robotic control that enables cooperative, tactically correct robotic behaviors that human teammates understand. For maximum effectiveness, unmanned systems (UMSs) must be able to support dismounted warfighters in high-intensity, high-operational-tempo (OPTEMPO) situations without becoming a source of distraction. Current models of robotic control require overt human tasking, limiting robotics to low OPTEMPO tasks. The Combat Causal Reasoner (CCR) proposes to change the paradigm of UMS autonomy by enabling UMSs to cooperate with humans without expecting the UMS to perceive the environment as a human would. CCR uses a Playbook approach to generate responses that are consistent with warfighter actions. An experiment demonstrated that a CCR-enabled robot measurably increased warfighter effectiveness and resource utilization, with no loss of robot effectiveness when compared to human tele-operation during high-tempo operations.

INTRODUCTION

Unmanned systems (UMS) on the battlefield provide warfighters with enhanced capabilities. UMSs are force multipliers that can assist troops to perform dangerous missions such as improvised explosive device (IED) defeat, perimeter security, and reconnaissance. With them, warfighters are more effective and more likely to survive the battlefield environment. For maximum effectiveness, UMSs must support warfighters in a wide variety of situations—especially in high-intensity, high-operational-tempo (OPTEMPO) situations—without becoming distractions. However, we must overcome the problem of increased end user cognitive burden that UMS technologies present.

Conventional models of robotic control are based on overt tasking from a human operator. Current command and control (C2) of UMSs fall into two main categories: tele-operation and map-based path planning. The inadequacies of these approaches are clear; both methods are too slow and too demanding for complex or chaotic military operations. Tele-operation is limited to tasks that do not require high OPTEMPO (e.g., IED disposal). Map-based methods require complex operator control units (OCU), high levels of autonomy, and the operator’s full attention. When the tactical situation changes, the operator must revise the robotic plan. Because of the complexity and limitations of map-based control systems, no UMS has been fielded with such a system.

Both tele-operation and map-based path planning are centered on OCU usage. OCUs are bulky, require specialized training, and consume too much of the operator’s attention. The most dangerous consequences of the OCU-based approach are the cognitive demands and attention drain on enmeshed warfighters. OCU-based control removes the unmanned system operator from the battle as an effective warfighter. In the USMC Urban Warrior exercises, the introduction of computer-based situational awareness systems actually *increased casualties* because the warfighters were too focused on their displays (Freedman, 2002).

Even with the levels of autonomy achieved in programs such as DARPA’s PerceptOR and Future Combat Systems’ ANS, the burden of tasking the UMS continues to fall to a human. UMSs cannot set their own objectives. Unfortunately, in military scenarios, the times and situations where robotic help is most needed occur when warfighters have the least time and attention to spare. Current robotic control requires the operator to go “heads-down” during robotic operations, situational awareness (SA) is significantly compromised, leaving him far more vulnerable. The few seconds that it takes to switch modes—drop the robotic hardware, grab a weapon, assess the situation, and take appropriate action—are seconds that the warfighter does not have. Warfighters in combat have too many tasks to dedicate much time to robotic control. UMSs must perform the right tasks at the right time, without significant attention from the warfighter.

The work described here will change the paradigm of UMS autonomy by enabling UMSs to cooperate with human warfighters. The approach does not expect the UMS to perceive the combat environment in the same way a human does. Combat Causal Reasoning (CCR) enables UMSs to recognize the actions and objectives of the warfighters and *generate its own support response* in real time. CCR reasons about the underlying goals of human activities and selects and executes the appropriate robotic behaviors. For a UMS to be truly effective in tactical situations, it must operate cooperatively with, and in close proximity to, human beings and human-occupied vehicles. The UMS must be able to recognize human activities in terms of root causes and tactical goals. It can then select and execute appropriate behaviors. The SAIC ACTR IRAD, which is the basis for our CCR effort, demonstrated that a UMS can effectively maneuver in coordination with a small unit by observing the motions of the unit members and executing complementary motions.

CCR software was developed and demonstrated in an operational scenario with a ground UMS. The CCR software was evaluated on its ability to recognize human behaviors and goals and to choose an appropriate complementary action, which it then executed. The CCR effort showed a measurable increase in warfighter effectiveness, resource utilization, and force multiplication.

SYSTEM DESCRIPTION

The CCR system enables UMSs to operate as integrated members of dismounted units. With CCR, UMSs develop cognitive models of individual warfighters and the unit as an organism. These cognitive models are used to interpret warfighter actions. The goal of CCR is *not* the complete understanding of the complexities of the tactical situation. Rather, the CCR identifies the commander's intent (CI), which enables the UMS to select and execute tactically relevant behavior without an OCU. The CCR frees the warfighter from continuous control duties, eliminates excess cognitive burden, and enables the UMSs to operate with little or no human attention during high-intensity operations. The result is a fully integrated team of warfighters and UMSs that work intuitively together without detailed plans, without keyboards and joysticks, and without constant overt control.

The key to successfully integrating the UMS with the team is to enable it to perform cooperative, tactically-correct behaviors without a complete situational understanding. In our concept, the UMS focuses on low-level, tactically-correct behaviors (as defined by Army FM 7-8) rather than complex interpretation. The extremely complex task of interpreting the environment and tactical situation is left to the warfighter.

With CCR, the warfighters' actions and reactions are cues that enable the UMS to select and perform appropriate actions. In effect, the warfighters become sensors and evaluators for the UMS. The UMS performs actions consistent with the warfighter actions; that is, the UMS runs when the warfighters run, hides when the warfighters hide, and performs learned and programmed tasks when the opportunity arises. Less simplistically, the UMS uses its warfighter/small-unit cognitive model to hypothesize tactical goals and the best course of action. Figure 1 illustrates the information flow through the CCR system.

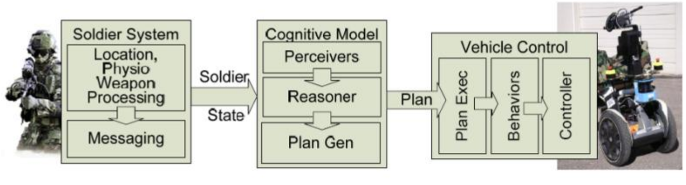

Figure 1. Information flow though the CCR system.

Perceivers

CCR Perceivers (see Table 1) are small software programs that acquire information from the squad and the environment that the CCR needs to make decisions. For example, based on warfighter position and weapons pointing, the Target Perceiver identifies possible targets. The possible targets are identified and a confidence value is assigned, reducing the reasoning challenge to identifying the appropriate behavior given the location and number of targets.

Table 1. Perceivers

Perceiver	Behavior	Sensors
Target Perceiver	Determine the location of hostile targets based on where BLUEFOR (friendly forces) warfighters are aiming their weapons	Warfighter Weapon Trigger sensors, Digital Compasses, and GPS
Gunfire Detection Perceiver	Determine if an enemy is firing near the BLUEFOR team location	Acoustic sensors (simulated)
Warfighter Weapon Perceiver	Provided data including how much ammunition remains	Weapons sensors that recorded last reload and number of shot fired
Warfighter Health Perceiver	Determines physical and mental state of the warfighter	Bioharness that recorded heart rate variability
Voice Perceiver	Enables the warfighter to override CCR Reasoner and select the "play"	Microphone input into netted communications
Formation Perceiver	Detects when warfighters are in a particular formation	GPS

Reasoner

The Reasoner must generate, in real-time, an appropriate robot response commensurate with the ongoing tasks and goals of the human squad, with minimal operator control. The mission support response must be relevant, of tactical utility, and not reduce lethality or survivability.

The Reasoner uses a Playbook approach (Miller, Pelican, & Goldman, 1999; Whitlow, Dorneich, Funk, & Miller, 2002). The key insight of the Playbook approach is that humans want to—and should be—in charge of operations, even when they cannot control every detail. In our vision, the UMS is a valuable and intelligent subordinate, but its value to the team decreases if more control is necessary to make it behave in tactically useful ways. The Playbook concept for coordinating activities of intelligent subordinates is inspired by American football. Team interactions in football are successful because:

- Players share a model of the domain, goals, tasks, and methods by which the goals can be accomplished.
- Players are intelligent and can operate with partial instructions, to interpret and adapt them to circumstances at run time, within the context of the CI.
- The football playbook permits changes to the play via activity sequence labeling that can be adapted as needed.

The CCR approach to C2 of UMS shares similar traits. Warfighters and the UMS share a domain model. The UMS is semi-intelligent and predictable. CCR reasons over a finite set of goals, plays, and behaviors. The predictability of behavior enables its actions to be appropriate to the situation. The UMS knows when no tactical actions are necessary, and reverts to its *default* behavior (e.g., follow). The strength of the CCR approach is that it does not attempt to be an autonomous, intelligent agent. The UMS is a useful, predictable asset that does what it knows how to do with minimal guidance.

A play is a generalized template for a short-term plan of action that is appropriate for the situation, well-understood by the squad members, and consistent with the expectations and training of the team. A play is expressed as the top layer of a hierarchical task network (HTN) in the language of tactics and doctrine. The Executable Plan Generator instantiates the play into a plan of executable UMS behaviors that realize the high-level goals specified by the play.

Play selection components allow the UMS to react to a situation with useful behaviors. Play selection is a process of reasoning over the team's and UMS robotic goals to determine actions most likely to support warfighters' tactics. The CCR maintains a set of all feasible plays for the current situation. The actions are specified by a play from a Playbook.

The Playbook approach provides control at a high level of abstraction, in a language that is both machine and human interpretable. The power of this approach is the shared understanding between the warfighters and the UMS. The terse interactions afforded by the Playbook approach allow rapid access to major plays and their alternatives—critical for effective robotic C2 in high OPTEMP situations.

We developed six plays for this program. The default play, 'Follow,' is what the UMS does if none of the conditions for the other Plays are feasible. Plays are assigned a tactical priority to resolve a situation where the conditions for more than one play are true (see Table 2).

Table 2. Plays implemented in CCR

Play Name	Rank	Conditions	Behavior
Follow	1	No other play active	UMS follows approx. 5m behind MG (machine gunner)
Attack Posture	2	Enemy detected, at least 2 BLUEFOR weapons off safe	UMS moved to a position 10m away from squad formation, 270 deg. from MG
Fire Support	3	Attack Posture condition plus recent friendly firing	UMS moved to a position 10m away from squad formation, 270 degrees from MG. UMS fires on OPFOR position as calculated by BLUFOR weapon vectors.
Resupply Ammo (to warfighter)	4	Recent friendly firing plus warfighter out of ammo (software counter)	UMS moved from current position to position directly behind the warfighter w/o ammo
Resupply Ammo (for UMS)	5	Recent friendly firing, ORP location specified, and UMS is out of ammo	UMS moved from current position to the ORP to receive additional ammo
Support (incapacitated warfighter)	6	Recent friendly firing, enemy detected, at least one warfighter NOT firing and exhibiting elevated stress levels	UMS moved from current position to a position between incapacitated warfighter and OPFOR—and maintained that position until firing behavior or stress level changed

METHOD

A field test evaluated the hypothesis that a CCR-enabled UMS, integrated with a small infantry unit, can have a substantial positive impact on attention and workload while remaining operationally effective.

Participants

Five participants formed the BLUEFOR (friendly force) unit. One participant played the role of robotics non-commissioned officer (RNCO), tasked with joysticking a robot in tele-operation trials and monitoring the robot during CCR trials to understand what activities it was executing. A squad leader (SL) led BLUFOR in tactical drills. The team leader (TL) and fire team member (FTM) carried out SL commands. The machine gunner (MG) carried a heavy weapon. The robot was assigned to assist the MG including following, re-supplying ammo, and flanking to provide additional fire on the enemy position. A single OPFOR (opposing force) provided resistance and opposition to BLUEFOR and was controlled by the experiment observer/controller (O/C) to achieve experimental objectives.

Our participants included veterans of Operation Iraqi Freedom. Our squad leader was a combat veteran M2/M3 Bradley platoon leader and our squad members included one battle-tested NCO. However, not all squad members had infantry experience, so basic training on maneuvering and rules of engagement were provided.

Experimental Design

The experiment utilized one independent variable, UMS control, with two levels: 1) CCR-enabled control (The UMS chooses behavior based on perceiver data. The human monitors UMS actions as needed) and 2) Tele-operation control (human operator selects UMS plays to execute based on situational assessment and executes them manually with a joystick controller).

Participants were asked to execute multiple isomorphic trials under both levels of the independent variable control condition. The order was not strictly counterbalanced.

Dependent Variables

Our performance metrics were:

- *Attention*. The percentage of time taken by the RNCO to command and control the UMS.
- *Workload*. We distributed the NASA TLX subjective workload survey (Hart & Staveland, 1988) to all BLUFOR participants at the end of each completed scenario. Ratings along the six scales ranged from 0-10
- *UMS effectiveness*. A custom instrument, the UMS Experience Survey, gauged participant impressions of the impact of the UMS in the tactical environment. Participants rated their level of agreement (1–7, strongly disagree to strongly agree) with six statements (described in Results).
- *Network Bandwidth*. This measure compared the amount of data required to control the UMS. We wanted to ensure that the move from tele-operation to a sensor-based CCR control mode did not increase network utilization, which could affect the deployability of the CCR approach. Network monitoring software was installed on the CCR vehicle's computer. As all network traffic was either destined to or

originated from the CCR vehicle, this approach accounted for all network traffic.

Tasks

The squad conducted an FM 7-8 style deliberate assault against a fixed position. We used the FM 7-8 task, conditions, and standards to measure tactical correctness. The CCR approach was compared to a tele-operation approach. For experimental purposes, the scenario was broken down into operationally meaningful chunks called vignettes. For all scenario runs in both control conditions, the robot was tasked with following the MG (default behavior), providing fire support upon contact with enemy, re-supply ammo when a squad member was out, and support incapacitated warfighter by the positioning the robot as a shield while the warfighter regained composure or un-jammed a weapon. In CCR-enabled mode, the robot was an autonomous, semi-intelligent squad member that could automatically conduct a limited repertoire of actions in response to sensed situations. In the tele-operated mode, the RNCO directed the robot to accomplish its tasks.

The RNCO monitored the UMS and maintained awareness of its behavior. In the tele-operation, the SA is a natural consequence of direct control of the robot. During CCR trials, the RNCO was instructed to raise his hand when actively monitoring the UMS; an experimenter incremented the cumulative time for the trials. The cumulative time was used to determine the percentage of each trial that the RNCO spent monitoring the UMS during CCR control condition. This time was compared to the percentage of time the RNCO tele-operated the UMS during—as calculated by the percentage of joysticking commands in network data logs.

Procedure

Participants conducted a deliberate attack on a fixed position in accordance with Army FM 7-8. Operations were conducted in an open area approximately 75 x 25 m. All participants received a drill briefing to familiarize them with the experimental activities. Participants also received training in the use of the custom Airsoft guns and associated safety precautions. The squad leader and RNCO received additional training in how to interact with the UMS and a detailed description of expected UMS behaviors during CCR trials.

BLUEFOR participants conducted a deliberate assault against an entrenched OPFOR, who was under the direct control of the O/C to insure the OPFOR behavior supported the experimental scenarios. BLUEFOR conducted an extended deliberate attack scenario with the UMS. The participant in this experiment was responsible for joysticking the UMS in tele-operation condition and monitoring UMS under CCR-enabled condition.

RESULTS

Attention

When manually controlling the vehicle via the OCU, the RNCO spent 70% of the total mission time controlling the vehicle. During CCR-enabled runs, an average of only 10% of the RNCO’s time was spent monitoring the vehicle. The result was significant (Two-Sample Student’s 2-tail t-test, t-value=7.0 and p=0.002). Note that it appears that as the participant gained more comfort with the vehicle in the CCR-enabled runs, vehicle monitoring time decreased greatly. More data points are necessary, but ignoring the first two CCR-enabled runs yields an average time devoted of only 2.9 %.

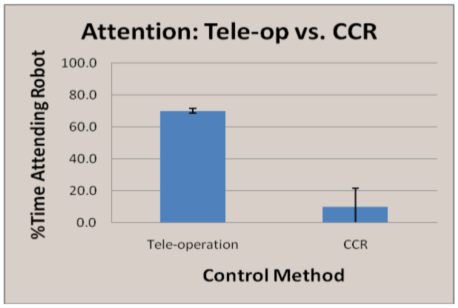

Figure 2. Percentage of time attending UMS.

Workload

We found no difference in perceived workload for all participants, validating the experimental assumption that there was nothing systematically different between the control conditions that impacted perceived workload of team members. Figure 3 illustrates the NASA TLX for the RNCO only. The RNCO reported substantially higher mental workload and temporal workload during tele-operation trials compared to CCR trials. This outcome was expected, given the all-encompassing nature of tele-operating the UMS.

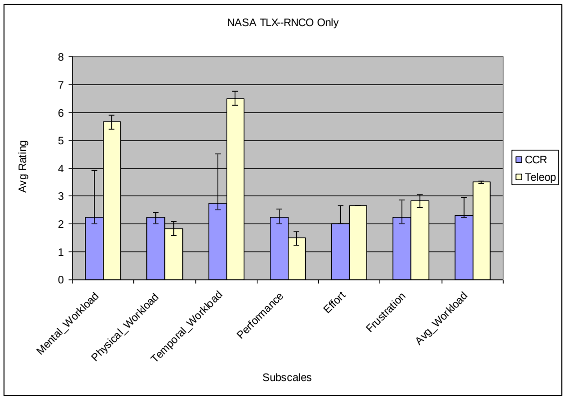

Figure 3. NASA TLX RNCO Only

Network Bandwidth

In tele-operation mode, the only recorded network traffic was that coming from the tele-operating warfighter system. Other network traffic, such as the Software Emergency Stop that was used both when operating the system in tele-operational mode and in CCR mode, was ignored, as was any auxiliary network traffic generated by engineering and monitoring tools. On average, the tele-operating system transmitted 483,935 bytes per second (BPS) during the experiment.

To calculate network utilization when the vehicle was operating in CCR enabled mode, all the traffic generated from the warfighter-worn sensors was combined. The CCR system used this data to determine the control of the vehicle. In CCR-enabled mode, the amount of network traffic required was limited to an average of 6,604 BPS. The difference was significant (Two-Sample Student's 2-tail t-test, t-value=58.9 and p=0.01)—a 73x decrease in bandwidth usage compared to tele-operating. Note that this traffic depends on the number of warfighters transmitting sensor data. As the number of warfighters increases, so will the required bandwidth. However, this linear growth will be so small (less than 2000 BPS per additional warfighter) that it is relatively insignificant.

UMS Experience Survey

Error! Reference source not found. shows that the control conditions did not produce significantly difference results. This validates our assumption that, in general, the experimental interaction with automation was not perceived to degrade the tactical performance of the squad.

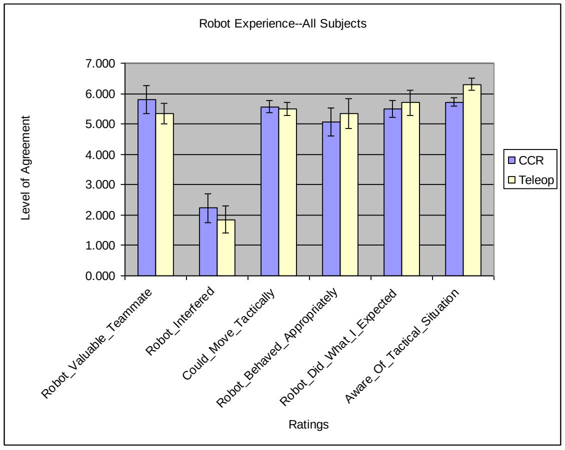

Figure 4. UMS experience survey. All participants.

Limitations

Given the preliminary nature of an early feasibility and field test, the evaluation presents several limitations. 1) The CCR could respond to only a limited action (play) repertoire and number of situations. In future evaluations, the playbook will be expanded. 2) The small sample of subjects, where some did not have small unit infantry operations, limited the statistical analysis possible. 3) The short length of the scenarios limited the data and time interacting with and observing the robot. 4) Limited realism of the scenarios made it difficult to produce the robust physiological stress level needed to consistently trigger the incapacitated-soldier play.

DISCUSSION

When the RNCO tele-operated the UMS, he was typically standing (i.e. not in a defensive posture), was not "in the fight" (i.e. engaging the enemy), and spent the majority of his attention on the UMS. In contrast, when the UMS was controlled by CCR, the RNCO was able to engage the enemy, take defensive postures that increased his safety, and generally support the primary goal of the mission. In addition, the RNCO's temporal and mental workload was lower during the CCR trials.

In both conditions, the members of the squad felt that the UMS was a valuable teammate, behaved appropriately, did what was expected, and was easy to control. Critically, no loss of these attributes were felt when UMS control was shifted from a dedicated human operator to the CCR approach. Thus, the results from this first experiment show that the CCR approach results in the squad gaining another warfighter with no loss of robotic asset effectiveness. Additional UMS warfighters become direct force multipliers. Their value-add is immediate and obvious.

Building on these results, research is required to gain a more detailed understating of the metrics and performance measures that can be used to quantify the impact the CCR-enabled UMS control paradigm could have on the battle space. Metrics should be developed with DARPA and potential transition partners in a series of increasingly relevant platforms and experimental vignettes. The range of plays should be expanded, as well as a deeper understanding of the triggers to initiate and to switch between plays.

Another promising area of research is the increased use of warfighter biometrics as inputs into the CCR. Longer term development of this technology would include research into the mixed-initiative aspects of CCR-enabled control through the use of voice overrides (the controller "calling a play"). We believe that UMS control should be primarily passive. When explicit control is needed, it should follow the natural form of communication that soldiers use with each other.

A follow-on effort should develop extensions of the voice grammar and develop a hand-signal grammar (through a data glove) to control the UMS as a natural extension of the CCR C2 technology. The addition of a learning component to the system would allow a CCR-enabled UMS to train with a squad to develop new plays or to tailor its behavior response to each squad's individual preferences. For example, one team may prefer that the UMS stay aggressive and on point while a different team may prefer to use the UMS as a cargo system and want it to stay a safe distance from the squad. Finally, research questions such as the transfer of authority, trust, and mode awareness all must be addressed if CCR is to go from promising concept to a robust control capability.

ACKNOWLEDGMENTS

This paper was supported by a contract with DARPA, NBCHC090016, for which Dr. Amy Kruse served as the Program Manager. The opinions expressed herein are those of the authors and do not necessarily reflect the views of DARPA.

REFERENCES

Freedman, D.H. (2002), “Killed at their Keyboards,” *Business 2.0*, February.

Hart, S.G., & Staveland, L.E. (1988). Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In P.A. Hancock & N. Meshkati (Eds.), *Human mental workload* (pp. 139–183). Amsterdam: North-Holland.

Miller, C., Pelican, M. and Goldman, R. (1999). “High Level ‘Tasking Interfaces’ for Uninhabited Combat Air Vehicles.” *In Proceedings of the International Conference on Intelligent User Interfaces*. January 5-8, Redondo Beach, CA.

Whitlow, S., Dorneich, M., Funk, H., and Miller, C. (2002). “Providing Appropriate Situation Awareness within a Mixed-Initiative Control System.” In *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, Hammamet, Tunisia; October 6-9.