Althea: Minimalist Representation for Robot Assembly Tasks

Jonathan B. Handler
Computer Science Department, Northwestern University
1890 Maple Avenue
Evanston, IL 60201
Email: jhandler@ils.nwu.edu; Phone: (847) 491-3500

Althea[1] is a hand-eye robot system that applies minimalist representation techniques (Agre and Chapman 1987; Brooks 1991; Horswill 1996) to the task of building a wall from Duplo bricks. Duplos are large Lego bricks — molded plastic blocks that interlock with one another. The system (figure 1) visually and then manually acquires Duplos, orients them in the robot's gripper, and then joins them to the wall.

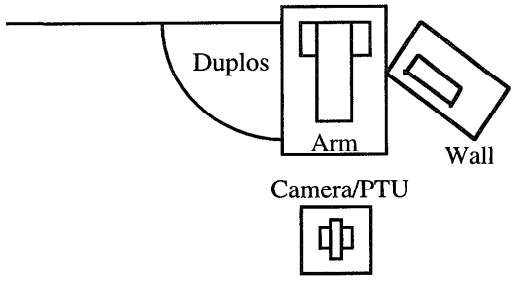

Figure 1: Top view of Althea's workspace.

Duplo acquisition is accomplished by servoing the camera to center the mass of Duplo pixels, yielding a pan/tilt coordinate. The system is provided with a table that translate these coordinates to arm coordinates. No attempt was made to cope with effects like foreshortening, and pose indeterminacy. Instead, the arm attempts to pick something up at each of the locations near the one corresponding to the final pan/tilt location. When something is in the gripper, then the system knows that it has succeeded.

Gripping the Duplo correctly is necessary before placement. The Duplo's orientation is changed relative to the gripper by a series of moves that place the Duplo on the table, reorient the hand and then regrip it. One way to pick the right reorientation routine to run is to calculate the Duplo's pose, then reason about adjacent faces. In the minimalist tradition, Althea instead calculates simple features of what is visible: is it a side? is its major visual axis horizontal or vertical? is there something that looks like the back of a Duplo? These features drive directly and robustly the choice of which rotation to perform.

The placement problem was restricted for Althea to cases in which there is at least one Duplo already on the wall. Althea is given the right edge of the wall as arm coordinates and moves the arm across the wall until it contacts the Duplo on the wall. It moves down, using the robot's wobble-mode until it contacts the wall, backs off and executes a high-speed ram into the newly-added piece.

Althea is an exercise in applying minimalist representation schemes to a domain that traditionally calls for much explicit representation. Although no rigorous testing was done, we can give a rough characterization of results. In acquisition, the system regularly acquires a Duplo without calculating pose, object type, or depth to object. Similarly, the system regularly reorients the Duplo correctly lacking a model of Duplo faces, accurate pose information, and a model of rotation. The placement routines were the least worked on, but even so, Althea was able to run end-to-end about half of the time. Handler (forthcoming) discusses these issues in greater detail.

When classical plans are executed in the natural world, they often fail. Knowing that the plan has failed can be difficult because the information sanctioning that inference — the envisioned world state when the plan was constructed — is not present in the plan. Minimalist representation schemes seek to ameliorate this problem through keeping less information.

Althea is a middle-ground alternative between classical planning and behavior-based systems that keeps just that information that is least likely to go out of date, while building into the system the knowledge it needs to update its model at appropriate times. By linking actions to their consequences, we increase the fidelity of inferences about the results of the actions. For instance a 66mm robot grip width may mean that there's a Duplo in the gripper. Then again, the gripper might just be stuck open. The raw grip width must be interpreted differently when the prior operation was a "close gripper" versus an "open gripper." Handler (forthcoming) discusses a plan executive that can use this kind of information.

Agre, P. E. and Chapman, D. (1987). "Pengi: an implementation of a theory of activity". In *Proceedings AAAI-87*, Seattle, Wa.

Brooks, R. (1991). "Intelligence without representation." Artificial Intelligence **47**: 139-160.

Handler, J. B. (forthcoming). Execution History as the Basis of Plan Execution in Autonomous Agents. Ph.D. thesis. Northwestern University, Computer Science.

Horswill, I. (1996). "How far can behavior-based architectures go?". In *SAB-96*, Cape Cod, MA, The MIT Press.

[1] See www.tinkertoy.ils.nwu.edu/jhandler/AAAI-97.ps.Z for a longer version of this paper.