COLLAGEN: travel assistant using shared graphical interface (Rich and Sidner)

- first design
  - agent and user interact with application, can communicate with each other, can observe what other has done with the application
  - behave according to principles of discourse: shared plan model
    * model segment purposes, shared goals
    * establishing mutual beliefs
    * producing and recognizing utterances
    * not NL: artificial language plus graphical manipulation

- example: user asks SF to Dallas to Atlanta for Wed 9am meeting
  - TT guides user through specification of constraints
  - TT uses domain-independent knowledge of human communication (Grosz and Sidner model of dialogue)
  - shared beliefs are required, yet not all in place at beginning
  - requires reasoning from partial knowledge

- principles underlying design
  - establish mutual belief for collaboration
  - distinguish shared and not, to communicate non-shared; need not deal with what is mutually believed
  - whenever detects user performing actions of possible collaboration: use communication to verify
  - TT awaits confirmation of things it does, to allow counterproposals
  - TT performs recipe recognition using observation of user's GUI actions and expects user intends to be recognized; checks recipe recognition with user
  - TT communicates re its own actions, if believed not yet shared

- idea of working together
  - start with partial knowledge, use communication, determine recipe
  - must share information
  - planning and acting interleaved
  - each segment extends the shared plan
  - agent can "wave" to signal attention
- usefulness of mutual belief
  - always check first with beliefs resulting from recipe recognition
  - assume certain things known (shared space)

- second model
  - all of the first model plus idea of having a history of the dialogue made available to the user
  - operations of return, replay or stop
    * retry: go back to old goal
    * revisit: pick up exactly where left off on old goal
    * stop: pop current segment off focus stack (forget goal)
    * undo: pretend you never did action
    * replay: reuse earlier work in slightly different way
  - support the user's problem solving process by relating current actions to the global context and history of the interaction
  - assume interface based on human discourse rules and conventions is easier for people to learn

- details on the model
  - agent observes actions of user, communicates with user, maintains history of collaboration, interacts directly with application
  - artificial language for negotiation represents the communication: e.g. offer, request, etc.
  - user can select "what have we been doing" from a communication menu
  - suggest what to do next by running discourse interpretation backwards, to get agenda of expected goals
  - interaction history becomes an object to manipulate