Designs for Phase I Trials of Combinations of Agents

Mark Conaway, University of Virginia
Nolan Wages, University of Virginia
John O'Quigley, Université de Paris VI
Alexia Iasonos, Memorial Sloan-Kettering

ENAR meeting 2012. Washington, DC

Outline

- Brief description of single agent phase I trials
- Examples of combination-agent trials
  - Enumerate all orderings
  - Many orderings
- Problems with small sample size and large number of possible “treatments” (combinations of agents)

Single Agent Phase I Trials

- Typical statistical set-up:
  - Preset dose levels $d_1 < d_2 < \ldots < d_K$
  - Binary measure of toxicity

  π_j = Prob patient receiving dose level j experiences a "dose-limiting toxicity" (DLT)

- Primary goal: Find maximum tolerated dose (MTD)

Single Agent Phase I trials

- MTD: highest dose that can be administered with an "acceptable" level of toxicity
  - "acceptable": Probability of toxicity is no more than a pre-specified amount
    - Often 20% or 33%
- Ethical considerations dictate that trials are done sequentially
  - Patients not allocated to dose level d_j unless levels $d_1, \ldots, d_{j-1}$ are believed to be "safe"

Many designs proposed in this setting

- Traditional (or “standard” or “3 + 3”)
- Storer 2-stage
- Up-and-down
- *Continual Reassessment Method (CRM)*
- Recently proposed Bayesian methods

CRM set-up

- Fixed number of dose levels: $d_1, d_2, \ldots, d_K$
- Use a "working model" for the probability of toxicity at dose level j:

$$\pi_j = (\psi_j)^a, \text{ where } 0 < \psi_1 < \psi_2 < \ldots \psi_K < 1$$

Ψ's are pre-set

'a' is a parameter to be estimated

Two-stage, likelihood version O'Quigley and Shen, Biocs 1996

- Stage I. Use any 'non-model' type design (any of Storer's stage 1, or up-and-down or..)
  - E.g. Start at dose level 1
  - Escalate in single patient cohorts
  - Once a toxicity is observed, start stage II
- Stage II: Have toxicity and number of patients {Yj, Nj} on dose levels 1, … , K
  - Likelihood:

$$\prod_{j=1}^{K} (\quad)^{y_j} (\quad)^{n_j - y_j}$$

Estimate ‘a’

- Estimate a by maximum likelihood $(\widehat{a})$
- Plug back into working model

$$\psi_1{}^{\hat{a}}, \psi_2{}^{\hat{a}}, \ldots, \psi_K{}^{\hat{a}}$$

- Next patient goes on dose level closest to target toxicity probability that defines the MTD

2-stage CRM

- Continue ‘estimate/allocate’ cycle until a fixed number of patients have been observed
- MTD estimate is recommended dose level for the next patient
- CRM has excellent statistical properties in terms of identifying the MTD

Partially-ordered trials

Combination	Pacitaxel	Carboplatin
1	54	6
2	67.5	6
3	81	6
4	94.5	6
5	67.5	7.5
6	67.5	9

Patnaik et al. (2000, Journal of Clin Onc)

- Toxicity probabilities follow a “partial order” : there exist pairs of combinations for which the ordering of toxicity probabilities is not known

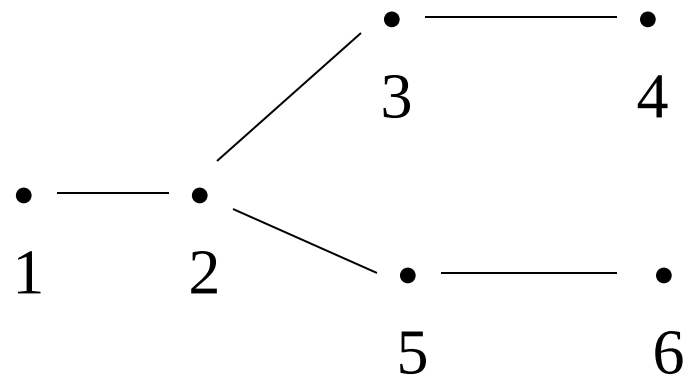

Compare to single agent trials

- Same:
  - Need to do the dose allocation sequentially

- Different
  - Toxicity probabilities follow a partial order

Wages, Conaway and O'Quigley (2011, Clinical Trials)

- Stage 1. Single patient escalation through "zones"

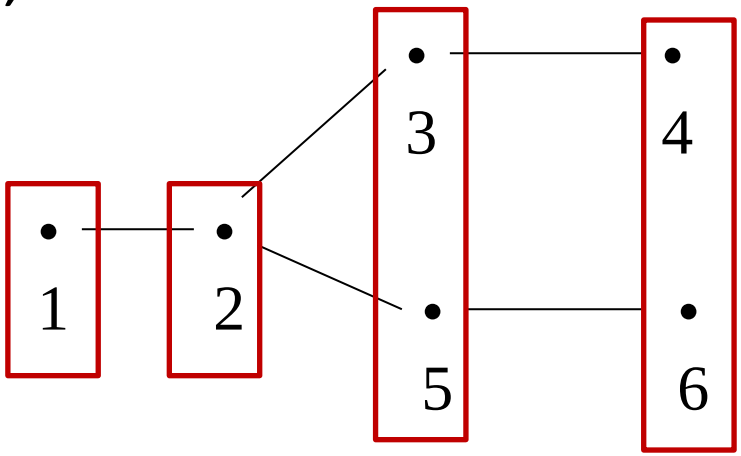

- After a toxicity is observed, start stage II.

CRM for Partial Orders

- Consider each (complete) order that is consistent with the partial order.
- Intuition: If we knew which one was the "correct" order, we could just use usual CRM

Comp Order	Ordering
M1	1 – 2 – 3 – 4 – 5 – 6
M2	1 – 2 – 3 – 5 – 4 – 6
M3	1 – 2 – 3 – 5 – 6 – 4
M4	1 – 2 – 5 – 3 – 4 – 6
M5	1 – 2 – 5 – 3 – 6 – 4
M6	1 – 2 – 5 – 6 – 3 – 4

CRM for partial orders

- 'Two-parameter' version of CRM
  - One parameter indexes the ordering
  - Within a given ordering, usual CRM set-up
- The working model for the probability of toxicity for combination i in ordering M=m is

$$\psi_{im}^{a_m}$$

Example of working model

M	1	2	3	4	5	6
1 (1-2-3-4-5-6)	$(.01)^{a1}$	$(.05)^{a1}$	$(.10)^{a1}$	$(.20)^{a1}$	$(.33)^{a1}$	$(.50)^{a1}$
2 (1-2-3-5-4-6)	$(.01)^{a2}$	$(.05)^{a2}$	$(.10)^{a2}$	$(.33)^{a2}$	$(.20)^{a2}$	$(.50)^{a2}$
…	…	…	…	…	…	..,
6 (1-2-5-6-3-4)	$(.01)^{a6}$	$(.05)^{a6}$	$(.33)^{a6}$	$(.50)^{a6}$	$(.10)^{a6}$	$(.20)^{a6}$

Working model consistent with the ordering.

Allocation method

- As data accumulates, estimate ‘a_m’ for each ordering by maximum likelihood
  - Choose ordering with largest likelihood
  - Update estimate of toxicity probabilities for dose combinations within that ordering
  - Next patient goes on dose combination with the estimated toxicity probability closest to the target

How well does it work?

- Wages, Conaway and O'Quigley (2011) present results of simulations assessing how well this identifies the MTD
- Comparisons to other methods for partially ordered trials:
  - Similar to Conaway, Dunbar and Peddada (2004) in identifying MTD
  - Not as often as CRM when you know the ordering

Illustration

Combination	True prob	Optimal Bench-mark	OS (1996) Correct Order	OS (1996) Incorrect Order	CDP (2004)	WCO (2011)
①	0.05	0	0	0	0.002	0
②	0.10	0.004	0.006	0.026	0.022	0.010
③	0.20	0.196	0.185	0.486	0.339	0.247
⑤	0.33	0.571	0.529	0.237	0.438	0.412
④	0.45	0.220	0.269	0.024	0.143	0.264
⑥	0.60	0.010	0.011	0.227	0.056	0.067
Summary		0.057	0.062	0.134	0.084	0.085

Summary : $\sum (\%\, recommended) \left| \pi_i - \text{target} \right|$

Without over-interpreting one set of true probabilities….

- If the ordering is known, problem reduces to single agent (usual) case
  - O'Quigley & Shen design gives results similar to optimal benchmark

- If guess incorrectly at the ordering and use a method relying on that ordering, poor properties in terms of estimating MTD

For one set of true probabilities…

- WCO and CDP have similar properties
  - Other cases, one may do better than the other, but in general similar properties
  - WCO computationally simpler
- Identifies MTD less often than when true ordering is known
  - In other cases, performance can be similar to case where true ordering is known

Could we weight the orderings?

- Wages, Conaway and O'Quigley (2011, Biometrics)
  - Uses model from first patient on
    - Not a 2-stage
    - Bayesian method
  - Allows prior weighting of orderings
  - Still considers all possible orderings consistent with partial order

Combination agent trials

		Dose of agent 2		
		0	1	2
Dose of agent 1	25	①	②	③
	100	④	⑤	⑥
	400	⑦	⑧	⑨
	1600	⑩	⑪	⑫

Methods

- Thall, Millikan, Mueller, Lee (2003, Biometrics)
- Conaway, Dunbar and Peddada (2004, Biometrics)
- Wang and Ivanova (2005, Biometrics)
- Yin and Yuan
  - 2008, Stat in Med
  - 2009 Applied Stat
  - 2009 Biometrics
- Braun and Wang (2010, Biometrics)
- Thall, Nguyen, Paoletti, Kramar (2010, Biometrics)
- Braun and Alonzo (2011, Clinical Trials)
- Wages, Conaway and O'Quigley
  - Biometrics, 2011
  - Clinical Trials, 2011
  - Under review, 2012

What makes this different?

- Stage I not different
  - Escalate through zones
  - Toxicity known to increase across zones, unknown within

	0	1	2
25	Zone 1	Zone 2	Zone 3
100	Zone 2	Zone 3	Zone 4
400	Zone 3	Zone 4	Zone 5
1600	Zone 4	Zone 5	Zone 6

Why is this different? Stage II

- Is it reasonable to consider all the possible orderings?
- If choose subset:
  - Is it important to have the correct order as one of the subset?
    - If yes, would that imply the subset should be large?
    - If no, would that imply the subset could be small?
      - Note: In the previous, the “correct” order was always in the set because we considered all of them

How to choose orderings?

- Type/dose of agents may give a ‘natural’ ordering
- Previous uses of these agents
- Spread them out over the design space
  - (J. Huesing)
- Choose ‘generic’ orders
  - Conjecture: these are sufficiently spread across the design space

Recommended set of orders

- Across columns

  1-2-3-4-5-6-7-8-9-10-11-12

- Down Rows:

  1-4-7-10-2-5-8-11-3-6-9-12

①	②	③
④	⑤	⑥
⑦	⑧	⑨
⑩	⑪	⑫

Recommended set of orders

- Diagonal ‘1’

  1-2-4-3-5-7-6-8-10-9-11-12

- Diagonal ‘1’ reversed within zones

  1-4-2-7-5-3-10-8-6-11-9-12

①	②	③
④	⑤	⑥
⑦	⑧	⑨
⑩	⑪	⑫

A couple more possibilities

- ‘Switchback 1’

  1-2-4-7-5-3-6-8-10-11-9-12

- ‘Switchback 2’

  1-4-2-3-5-7-10-8-6-9-11-12

①	②	③
④	⑤	⑥
⑦	⑧	⑨
⑩	⑪	⑫

What effect does the choice have?

- Wages, O'Quigley and Conaway (submitted) investigate a 4 x 4 case
- Consider the use of 3, 6, or 9 orders
- Answer is complicated: depends on where MTD is in the table

In general

- 6 chosen orders
  - provides a good compromise even when ‘true’ ordering is not one of the set.
  - At times, can perform nearly as well as knowing the ordering.

Summary

- Generalization of CRM to partial orders
  - Good properties when it is possible to enumerate all orderings
- When it is not possible to enumerate orderings
  - Can incorporate prior knowledge of orderings
  - Has good properties when ‘general’ choice of orderings is used.